

RESEARCH CENTRE

Rennes - Bretagne Atlantique

IN PARTNERSHIP WITH:

Ecole Nationale Supérieure
Mines-Télécom Atlantique Bretagne Pays
de la Loire

2021

ACTIVITY REPORT

Project-Team

STACK

Software Stack for Massively Geo-Distributed Infrastructures

IN COLLABORATION WITH: Laboratoire des Sciences du numérique de
Nantes

DOMAIN

Networks, Systems and Services,
Distributed Computing

THEME

Distributed Systems and middleware

Contents

| | |
|---|-----------|
| Project-Team STACK | 1 |
| 1 Team members, visitors, external collaborators | 2 |
| 2 Overall objectives | 3 |
| 2.1 STACK in a Nutshell | 3 |
| 2.2 Toward a STACK for Geo-Distributed Infrastructures | 3 |
| 3 Research program | 4 |
| 3.1 Overview | 4 |
| 3.2 Resource Management | 4 |
| 3.3 Programming Support | 6 |
| 3.4 Energy | 8 |
| 3.5 Security | 9 |
| 4 Application domains | 10 |
| 4.1 Overview | 10 |
| 4.2 Industrial Internet | 11 |
| 4.3 Internet of Skills | 11 |
| 4.4 e-Health | 11 |
| 4.5 Network Virtualization and Mobile Edge Services | 11 |
| 5 Social and environmental responsibility | 12 |
| 5.1 Footprint of research activities | 12 |
| 5.2 Impact of research results | 12 |
| 6 Highlights of the year | 12 |
| 6.1 Awards | 12 |
| 7 New software and platforms | 13 |
| 7.1 New software | 13 |
| 7.1.1 MAD | 13 |
| 7.1.2 Nitro | 14 |
| 7.1.3 VMPlaceS | 14 |
| 7.1.4 ENOS | 14 |
| 7.1.5 EnOSlib | 15 |
| 7.1.6 Concerto | 15 |
| 7.2 New platforms | 16 |
| 7.2.1 OpenStack | 16 |
| 7.2.2 Grid'5000 | 16 |
| 7.2.3 SeDuCe | 16 |
| 7.2.4 PiSeDuCe | 16 |
| 7.2.5 SILECS/SLICES | 16 |
| 8 New results | 17 |
| 8.1 Resource Management | 17 |
| 8.2 Programming Support | 19 |
| 8.3 Energy-aware computing | 20 |
| 8.4 Security and Privacy | 20 |
| 9 Bilateral contracts and grants with industry | 21 |
| 9.1 Bilateral contracts with industry | 21 |
| 9.2 Bilateral grants with industry | 22 |

| | |
|--|-----------|
| 10 Partnerships and cooperations | 22 |
| 10.1 International research visitors | 22 |
| 10.2 European initiatives | 23 |
| 10.2.1 FP7 & H2020 projects | 23 |
| 10.2.2 Other european programs/initiatives | 23 |
| 10.3 National initiatives | 23 |
| 10.3.1 Ademe | 23 |
| 10.3.2 CNRS | 24 |
| 10.3.3 ANR | 24 |
| 10.3.4 PIA 4 | 25 |
| 10.3.5 Etoiles Montantes | 25 |
| 10.4 Regional initiatives | 25 |
| 10.4.1 Interregional activities | 25 |
| 10.4.2 Nantes excellency initiative in Medecine and Informatics (NExT) | 26 |
| 11 Dissemination | 27 |
| 11.1 Promoting scientific activities | 27 |
| 11.1.1 Scientific events: organisation | 27 |
| 11.1.2 Scientific events: selection | 27 |
| 11.1.3 Journal | 27 |
| 11.1.4 Scientific expertise | 27 |
| 11.1.5 Research administration | 28 |
| 11.2 Teaching - Supervision - Juries | 28 |
| 11.2.1 Teaching | 28 |
| 11.2.2 Supervision | 29 |
| 11.2.3 Juries | 29 |
| 11.3 Popularization | 30 |
| 11.3.1 Articles and contents | 30 |
| 11.3.2 Education | 30 |
| 12 Scientific production | 30 |
| 12.1 Major publications | 30 |
| 12.2 Publications of the year | 30 |
| 12.3 Cited publications | 32 |

Project-Team STACK

Creation of the Project-Team: 2019 January 01

Keywords

Computer sciences and digital sciences

- A1.1.8. – Security of architectures
- A1.1.10. – Reconfigurable architectures
- A1.1.13. – Virtualization
- A1.3.4. – Peer to peer
- A1.3.5. – Cloud
- A1.3.6. – Fog, Edge
- A1.5.1. – Systems of systems
- A1.6. – Green Computing
- A2.1.7. – Distributed programming
- A2.1.10. – Domain-specific languages
- A2.5.2. – Component-based Design
- A2.6. – Infrastructure software
- A2.6.1. – Operating systems
- A2.6.2. – Middleware
- A2.6.3. – Virtual machines
- A2.6.4. – Ressource management
- A3.1.2. – Data management, quering and storage
- A3.1.3. – Distributed data
- A3.1.8. – Big data (production, storage, transfer)
- A4.1. – Threat analysis
- A4.4. – Security of equipment and software
- A4.9. – Security supervision

Other research topics and application domains

- B2. – Health
- B4. – Energy
- B4.5.1. – Green computing
- B5.1. – Factory of the future
- B6.3. – Network functions
- B6.4. – Internet of things
- B6.5. – Information systems
- B7. – Transport and logistics
- B8. – Smart Cities and Territories

1 Team members, visitors, external collaborators

Faculty Members

- Adrien Lebre [Team leader, IMT Atlantique, Professor, HDR]
- H el ene Coullon [IMT Atlantique, Chair]
- Remous Aris Koutsiamanis [IMT Atlantique, Associate Professor]
- Thomas Ledoux [IMT Atlantique, Professor, HDR]
- Jean-Marc Menaud [IMT Atlantique, Professor, HDR]
- Jacques Noy e [IMT Atlantique, Associate Professor]
- Mario S udholt [IMT Atlantique, Professor, HDR]

Post-Doctoral Fellow

- Abdelghani Alidra [IMT Atlantique, from May 2021]

PhD Students

- Geo Johns Antony [Inria, from Apr 2021]
- Hiba Awad [Alterway, CIFRE, from Nov 2021]
- Maxime Belair [Orange Labs, CIFRE, until Sep 2021]
- Marie Delavergne [Inria]
- David Espinel [Orange Labs, CIFRE, until Apr 2021]
- Wilmer Edicson Garzon Alfonso [IMT Atlantique, until Jul 2021]
- Pierre Jacquet [Inria, from Oct 2021]
- Antoine Omond [IMT Atlantique, from Dec 2021]
- Jolan Philippe [IMT Atlantique]
- Dimitri Saingre [Armines, until Nov 2021]
- Sirine Sayadi [IMT Atlantique]

Technical Staff

- Ronan-Alexandre Cherrueau [Inria, Engineer, until Jun 2021]
- Brice Nedelec [Sigma / Inria, Engineer, from Jan 2021 to Mar 2021 and from Sept 2021 to Dec 2021]
- Remy Pottier [LS2N, Engineer, until Oct 2021]
- Simon Robillard [IMT Atlantique, Engineer, until Jun 2021]

Interns and Apprentices

- Matthieu Juzdzewski [Inria, from Feb 2021 until Jun 2021]
- Arnaud Szymanek [Inria, from Feb 2021 until Aug 2021]

Administrative Assistant

- Anne-Claire Binétruy [Inria]

2 Overall objectives

2.1 STACK in a Nutshell

The STACK team addresses challenges related to the management and advanced usages of Utility Computing infrastructures (i.e., Cloud, Fog, Edge, and beyond). More specifically, the team is interested in delivering appropriate system abstractions to operate and use massively geo-distributed infrastructures, from the lowest (system) levels to the highest (application development) ones, and addressing cross-cutting dimensions such as energy or security.¹ These infrastructures are critical for the emergence of new kinds of applications related to the digitalization of the industry and the public sector (a.k.a. the Industrial and Tactile Internet).

2.2 Toward a STACK for Geo-Distributed Infrastructures

With the advent of Cloud Computing, modern applications have been developed on top of advanced software stacks composed of low-level system mechanisms, advanced middleware and software abstractions. While each of these layers has been designed to enable developers to efficiently use ICT resources without dealing with the burden of the underlying infrastructure aspects, the complexity of the resulting software stack has become a key challenge. As an example, Map/Reduce frameworks such as Hadoop have been developed to benefit from the cpu/storage capacities of distinct servers. Running such frameworks on top of a virtualized cluster (i.e., in a Cloud) can lead to critical situations if the resource management system decides to consolidate all the VMs on the same physical machine [95]. In other words, self-management decisions taken in isolation at one level (infrastructure, middleware, or application) may indirectly interfere with the decisions taken by another layer, and globally affect the performance of the whole stack. Considering that geo-distributed ICT infrastructures significantly differ from the Cloud Computing ones regarding heterogeneity, resiliency, and the potential massive distribution of resources and networking environments [56, 89], we can expect that the complexity of the software stacks is going to increase. Such an assumption can be illustrated, for instance, by the software architecture proposed in 2016 by the ETSI Mobile edge computing Industry Specification Group [81]. This architecture is structured around a new layer in charge of orchestrating distinct independent cloud systems, *a.k.a.* Virtual Infrastructure Managers (VIMs) in their terminology. By reusing VIMs, ETSI targets an edge computing resource management that behaves in the same fashion as Cloud Computing ones. While mitigating development requirements, such a proposal hides all the management decisions that might be taken in the VIM of one particular site and thus may lead to conflicting decisions and consequently to non-desired states overall.

Through the STACK team, we propose to investigate the software stack challenge as a whole. We claim it is the only way to limit as much as possible the complexity of the next generation software stack of geo-distributed ICT infrastructures. To reach our goal, we will identify major building blocks that should compose such a software stack, how they should be designed (i.e., from the internal algorithms to the APIs they should expose), and finally how they should interact with each other.

Delivering such a software stack is an ambitious objective that goes beyond the activities of one research group. However, our expertise, our involvements in different consortiums (such as OpenStack) as well as our participation in different collaborative projects enable STACK members to contribute to this challenge in terms of architecture models, distributed system mechanisms and software artefacts, and finally, guideline reports on opportunities and constraints of geo-distributed ICT infrastructures.

¹Following the recruitment of Remous-Aris Koutsiamanis in October 2020, the team broadened its activities to the IoT area, covering now the continuum of Cloud-IoT devices under the resource management perspective.

3 Research program

3.1 Overview

STACK research activities have been organized around four research topics. The two first ones are related to the resource management mechanisms and the programming support that are mandatory to operate and use ICT geo-distributed resources (compute, storage, network, and IoT devices). They are transverse to the System/Middleware/Application layers, which generally compose a software stack, and nurture each other (*i.e.*, the resource management mechanisms will leverage abstractions/concepts proposed by the programming support axis and reciprocally). The third and fourth research topics are related to the Energy and Security dimensions (both also crosscutting the three software layers). Although they could have been merged with the first two axes, we identified them as independent research directions due to their critical aspects with respect to the societal challenges they represent. In the following, we detail the actions we plan to do in each research direction.

3.2 Resource Management

The challenge in this axis is to identify, design or revise mechanisms that are mandatory to operate and use a set of massively geo-distributed resources in an efficient manner [43]. This covers considering challenges at the scale of nodes, within one site (*i.e.*, one geographical location) and throughout the whole geo-distributed ICT infrastructure. It is noteworthy that the network community has been investigating similar challenges for the last few years [62]. To benefit from their expertise, in particular on how to deal with intermittent networks, STACK members have recently initiated exchanges and collaborative actions with some network research groups and telcos (see Section 10). We emphasize, however, that we do not deliver contributions related to network equipments/protocols. The scientific and technical achievements we aim to deliver are related to the (distributed) system aspects.

Performance Characterization of Low-Level Building Blocks Although Cloud Computing has enabled the consolidation of services and applications into a subset of servers, current operating system mechanisms do not provide appropriate abstractions to prevent (or at least control) the performance degradation that occurs when several workloads compete for the same resources [95]. Keeping in mind that server density is going to increase with physical machines composed of more and more cores and that applications will be more and more data intensive, it is mandatory to identify interferences that appear at a low level on each dimension (compute, memory, network, and storage) and propose counter-measures. In particular, previous studies [95, 54] on pros and cons of current technologies – virtual machines (VMs) [68, 76], containers and microservices – which are used to consolidate applications on the same server, should be extended: In addition to evaluating the performance we can expect from each of these technologies on a single node, it is important to investigate interferences that may result from cross-layer and remote communications [96]. We will consider in particular all interactions related to geo-distributed systems mechanisms/services that are mandatory to operate and use geo-distributed ICT infrastructures.

Geo-Distributed System Mechanisms Although several studies have been highlighting the advantages of geo-distributed ICT infrastructures in various domains (see Section 4), progress on how to operate and use such infrastructures is marginal. Current solutions [25] [26] are rather close to the initial Cisco Fog Computing proposal that only allows running domain-specific applications on edge resources and centralized Cloud platforms [33] (in other words, these solutions do not allow running stateful workloads in isolated environments such as containers or VMs). More recently, solutions leveraging the idea of federating VIMs (as the aforementioned ETSI MEC proposal [81]) have been proposed. ONAP [88], an industry-driven solution, enables the orchestration and automation of virtual network functions across distinct VIMs. From the academic side, FogBow [36] aims to support federations of Infrastructure-as-a-Service (IaaS) providers. Finally, NIST initiated a collaborative effort with IEEE to advance Federated Cloud platforms through the development of a conceptual architecture and a vocabulary². Although all these projects provide valuable contributions, they face the aforementioned orchestration limitations

²NIST and IEEE Federated Cloud conceptual architecture and vocabulary draft (Dec 2018).

(*i.e.*, they do not manage decisions taken in each VIM). Moreover, they all have been designed by only considering the developer/user's perspective. They provide abstractions to manage the life cycle of geo-distributed applications, but do not deliver means to administer the physical resources.

To cope with specifics of Wide-Area networks while delivering most features that made Cloud Computing solutions successful also at the edge, our community should first identify limitations/drawbacks of current resource management system mechanisms with respect to the Fog/Edge requirements and propose revisions when needed [61, 74].

To achieve this aim, STACK members propose to conduct first a series of studies aiming at understanding the software architecture and footprint of major services that are mandatory for operating and using Fog/Edge infrastructures (storage backends, monitoring services, deployment/reconfiguration mechanisms, etc.). Leveraging these studies, we will investigate how these services should be deployed in order to deal with resources constraints, performance variability, and network split brains. We will rely on contributions that have been accomplished in distributed algorithms and self-* approach for the last decade. In the short and medium term, we plan to evaluate the relevance of NewSQL systems [46] to store internal states of distributed system mechanisms in an edge context, and extend our proposals on new storage backends such as key/value stores [45, 87], and burst buffers [97]. We also plan to conduct new investigations on data-stream frameworks for Fog and Edge infrastructures [40]. These initial contributions should enable us to identify general rules to deliver other advanced system mechanisms that will be mandatory at the higher levels in particular for the deployment and reconfiguration manager in charge of orchestrating all resources.

Capacity Planning and Placement Strategies An objective shared by users and providers of ICT infrastructures is to limit as much as possible the operational costs while providing the expected and requested quality of service (QoS). To optimize this cost while meeting QoS requirements, data and applications have to be placed in the best possible way onto physical resources according to data sources, data types (stream, graphs), application constraints (real-time requirements) and objective functions. Furthermore, the placement of applications must evolve through time to cope with the fluctuations in terms of application resource needs as well as the physical events that occur at the infrastructure level (resource creation/removals, hardware failures, etc.). This placement problem, *a.k.a.* the deployment and reconfiguration challenge as it will be described in Section 3.3, can be modeled in many different ways, most of the time by multi-dimensional and multi-objective bin-packing problems or by scheduling problems which are known to be difficult to solve. Many studies have been performed, for example, to optimize the placement of virtual machines onto ICT infrastructures [70]. STACK will inherit the knowledge acquired through previous activities in this domain, particularly its use of constraint programming strategies in autonomic managers [66, 65], relying on MAPE (monitor, analyze, plan, and execute) control loops. While constraint programming approaches are known to hardly scale, they enable the composition of various constraints without requiring to change heuristic algorithms each time a new constraint has to be considered [64]. We believe it is a strong advantage to deal with the diversity brought by geo-distributed ICT infrastructures. Moreover, we have shown in previous work that decentralized approaches can tackle the scalability issue while delivering placement decisions good enough and sometimes close to the optimal [80].

Leveraging this expertise, we propose, first, to identify new constraints raised by massively geo-distributed infrastructures (*e.g.*, data locality, energy, security, reliability and the heterogeneity and mobility of the underlying infrastructure). Based on this preliminary study, we will explore new placement strategies not only for computation sandboxes but for data (location, replication, streams, etc.) in order to benefit from the geo-distribution of resources and meet the required QoS. These investigations should lead to collaborations with operational research and optimization groups such as TASC, another research group from IMT Atlantique.

Second, we will leverage contributions made on the previous axis "Performance Characterization of Low-Level Building Blocks" to determine how the deployment of the different units (software components and data sets) should be executed in order to reduce as much as possible the time to reconfigure the system (*i.e.*, the *Execution* phase in the control loop). In some recent work [76], we have shown that the provisioning of a new virtual machine should be done carefully to mitigate boot penalties. More generally, proposing an efficient action plan for the *Execution* phase will be a major point as Wide-Area-Network specifics may lead to significant delays, in particular when the amount of data to be manipulated is

important.

Finally, we will investigate new approaches to decentralize the placement process while considering the geo-distributed context. Among the different challenges to address, we will study how a combination of autonomic managers, at both the infrastructure and application levels [53], could be proposed in a decentralized manner. Our first idea is to geo-distribute a fleet of small control loops over the whole infrastructure. By improving the locality of data collection and weakening combinatorics, these loops would allow the system to address responsiveness and quality expectations.

3.3 Programming Support

We pursue two main research directions relative to new programming support: first, developing new programming models with appropriate support in existing languages (libraries, embedded DSLs, etc.) and, second, providing new means for deployment and reconfiguration in geo-distributed ICT environments, principally supporting the mapping of software onto the infrastructure. For both directions two levels of challenges are considered. On the one hand, the *generic* level refers to efforts on programming support that can be applied to any kind of distributed software, application or system. On this level, contributions could thus be applied to any of the three layers addressed by STACK (*i.e.*, system, middleware or application). On the other hand, the corresponding generic programming means may not be appropriate in practice (*e.g.*, requirements for more dedicated support, performance constraints, etc.), even if they may lead to interesting general properties. For this reason, a *specific* level is also considered. This level could be based on the generic one but addresses specific cases or domains.

Programming Models and Languages Extensions The current landscape of programming support for cloud applications is fragmented. This fragmentation is based on apparently different needs for various kinds of applications, in particular, web-based, computation-based, focusing on the organization of the computation, and data-based applications, within the last case a quite strong dichotomy between applications considering data as sets or relations, close to traditional database applications and applications considering data as real-time streams. This has led to various programming models, in a loose sense, including for instance microservices, graph processing, dataflows, streams, etc. These programming models have mostly been offered to the application programmer in the guise of frameworks, each offering subtle variants of the programming models with various implementation decisions favoring particular application and infrastructure settings. Whereas most frameworks are dedicated to a given programming model, *e.g.*, basic Pregel [75], Hive [91], Hadoop [92], some of them are more general-purpose through the provision of several programming models, *e.g.*, Flink [39] and Spark [72]. Finally, some dedicated language support has been considered for some models (*e.g.*, the language SPL underlying IBM Streams [67]) as well as core languages and calculi (*e.g.*, [35, 85]).

This situation raises a number of challenges on its own, related to a better structuring of the landscape. It is necessary to better understand the various programming models and their possible relations, with the aim of facilitating, if not their complete integration, at least their composition, at the conceptual level but also with respect to their implementations, as specific languages and frameworks.

Switching to massively geo-distributed infrastructures adds to these challenges by leading to a new range of applications (*e.g.*, smart-* applications) that, by nature, require mixing these various programming models, together with a much more dynamic management of their runtime.

In this context, STACK would like to explore two directions:

- First, we propose to contribute to generic programming models and languages to address composability of different programming models [48]. For example, providing a generic stream data processing model that can operate under both data stream [39] and operation stream [98] modes, thus streams can be processed in micro batches to favour high throughput or record by record to sustain low latency. Software engineering properties such as separation of concerns and composition should help address such challenges [28, 86]. They should also facilitate the software deployment and reconfiguration challenges discussed below.
- Second, we plan to revise relevant programming models, the associated specific languages, and their implementation according to the massive geo-distribution of the underlying infrastructure,

the data sources, and application end-users. For example, although SPL is extensible and distributed, it has been designed to run on multi-cores and clusters [67]. It does not provide the level of dynamicity required by geo-distributed applications (*e.g.*, to handle topology changes, loss of connectivity at the edge, etc.).

Moreover, as more network data transfers will happen within a massively geo-distributed infrastructure, correctness of data transfers should be guaranteed. This has potential impact from the programming models to their implementations.

Deployment and Reconfiguration Challenges The second research direction deals with the complexity of deploying distributed software (whatever the layer, application, middleware or system) onto an underlying infrastructure. As both the deployed pieces of software and the infrastructures addressed by STACK are large, massively distributed, heterogeneous and highly dynamic, the deployment process cannot be handled manually by developers or administrators. Furthermore, and as already mentioned in Section 3.2, the initial deployment of some distributed software will evolve through time because of the dynamicity of both the deployed software and the underlying infrastructures. When considering reconfiguration, which encompasses deployment as a specific case, the problem becomes more difficult for two main reasons: (1) the current state of both the deployed software and the infrastructure has to be taken into account when deciding on a reconfiguration plan, (2) as the software is already running the reconfiguration should minimize disruption time, while avoiding inconsistencies [73, 78]. Many deployment tools have been proposed both in academia and industry [50]. For example, **Ansible**, **Chef** and **Puppet** are very well-known generic tools to automate the deployment process through a set of batch instructions organized in groups (*e.g.*, *playbooks* in Ansible). Some tools are specific to a given environment, like Kolla to deploy OpenStack, or the embedded deployment manager within Spark. Few reconfiguration capabilities are available in production tools such as *scaling* and *restart* after a fault, for example in **Kubernetes** and **Juju charms**. Academia has contributed to generic deployment and reconfiguration models. Most of these contributions are component-based. Component models divide a distributed software as a set of component instances (or modules) and their assembly, where components are connected through well defined interfaces [86]. Thus, modeling the reconfiguration process consists in describing the life cycle of different components and their interactions. Most component-based approaches offer a fixed life cycle, *i.e.*, identical for any component [55]. Two main contributions are able to customize life cycles, Fractal [38, 31] and its evolutions [28, 29, 52], and Aeolus [47]. In Fractal, the *control* part of a component (*e.g.*, its life cycle) is modeled itself as a component assembly that is highly flexible. Aeolus, on the other hand, offers a finer control on both the evolution and the synchronization of the deployment process by modeling each component life cycle with a finite state machine.

A reconfiguration raises at least five questions, all of them are correlated: (1) *why software has to be reconfigured?* (monitoring, modeling and analysis) (2) *what should be reconfigured?* (software modeling and analysis), (3) *how should it be reconfigured?* (software modeling and planning decisions), (4) *where should it be reconfigured?* (infrastructure modeling and planning decisions), and (5) *when to reconfigure it?* (scheduling algorithms). STACK will contribute to all aspects of a reconfiguration process as described above. However, according to the expertise of STACK members, we will focus mainly on the three first questions: *why*, *what* and *how*, leaving questions *where* and *when* to collaborations with operational research and optimization teams.

First of all, we would like to investigate *why software has to be reconfigured?* Many reasons could be mentioned, such as hardware or software fault tolerance, mobile users, dynamicity of software services, etc. All those reasons are related somehow to the Quality of Service (QoS) or the Service Level Agreement (SLA) between the user and the Cloud provider. We first would like to explore the specificities of QoS and SLAs in the case of massively geo-distributed ICT environments [82]. By being able to formalize this question, analyzing the requirement of a reconfiguration will be facilitated.

Second, we think that four important properties should be enhanced when deploying and reconfiguring models in massively geo-distributed ICT environments. First, as low-latency applications and systems will be subject to deployment and reconfiguration, the performance and the ability to scale are important. Second, as many different kinds of deployments and reconfigurations will concurrently hold within the infrastructure, processes have to be reliable, which is facilitated by a fine-grained control of the process. Finally, as many different software elements will be subject to deployment and reconfiguration,

common generic models and engines for deployment and reconfiguration should be designed [37]. For these reasons, we intend to go beyond Aeolus by: first, leveraging the expression of parallelism within the deployment process, which should lead to better performance; second, improving the separation of concerns between the component developer and the reconfiguration developer; third, enhancing the possibility to perform concurrent and decentralized reconfigurations.

Research challenges relative to programming support have been presented above. Many of these challenges are related, in different manners, to the resource management level of STACK or to crosscutting challenges, *i.e.*, energy and security. First, one can notice that any programming model or deployment and reconfiguration implementation should be based on mechanisms related to resource management challenges. For this reason, all challenges addressed within this section are linked with lower level building blocks presented in Section 3.2. Second, as detailed above, deployment and reconfiguration address at least five questions. The question *what?* is naturally related to programming support. However, questions *why*, *how?*, *where?* and *when?* are also related to Section 3.2, for example, to monitoring and capacity planning. Moreover, regarding the deployment and reconfiguration challenges, one can note that the same goals recursively happen when deploying the control building blocks themselves (bootstrap issue). This comforts the need to design generic deployment and reconfiguration models and frameworks. These low-level models should then be used as back-ends to higher-level solutions. Finally, as *energy* and *security* are crosscutting themes within the STACK project, many additional energy and security considerations could be added to the above challenges. For example, our deployment and reconfiguration frameworks and solutions could be used to guarantee the deployment of end-to-end security policies or to answer specific energy constraints [63] as detailed in the next section.

3.4 Energy

The overall electrical consumption of DCs grows according to the demand of Utility Computing. Considering that the latter has been continuously increasing since 2008, the energy footprint of Cloud services overall is nowadays critical with about 91 billion kilowatt-hours of electricity [84]. Besides the ecological impact, the energy consumption is a predominant criterion for providers since it determines a large part of the operational cost of their infrastructure. Among the different approaches that have been investigated to reduce the energy footprint, some studies have been investigating the use of renewable energy sources to power microDCs [57]. Workload distribution for geo-distributed DCs is also another promising approach [59, 71, 93]. Our research will extend these results with the ultimate goal of considering the different opportunities to control the energy footprint across the whole stack (hardware and software opportunities, renewable energy, thermal management, etc.). In particular, we identified several challenges that we will address in this context within the STACK framework.

First, we propose to evaluate the energy efficiency of low-level building blocks, from the viewpoints of computation (VMs, containers, microkernel, microservices) [51] and data (hard drives, SSD, in-memory storage, distributed file systems). For computations, in the continuity of our previous work [49, 66], we will investigate workload placement policies according to energy (minimizing energy consumption, power capping, thermal load balancing, etc.). Regarding the data dimension, we will investigate, in particular, the trade-offs between energy consumption and data availability, durability and consistency [44, 87]. Our ambition is to propose an adaptive energy-aware data layout and replication scheme to ensure data availability with minimal energy consumption. It is noteworthy that these new activities will also consider our previous work on DCs partially powered by renewable energy (see the SeDuCe project, in Section 7.2), with the ultimate goal of reducing the CO₂ footprint.

Second, we will complete current studies to understand pros and cons of massively geo-distributed infrastructures from the energy perspective. Addressing the energy challenge is a complex task that involves considering several dimensions such as the energy consumption due to the physical resources (CPU, memory, disk, network), the performance of the applications (from the computation and data viewpoints), and the thermal dissipation caused by air conditioning in each DC. Each of these aspects can be influenced by each level of the software stack (*i.e.*, low-level building blocks, coordination and autonomous loops, and finally application life cycle). In previous projects, we have studied and modeled the consumption of the main components, notably the network, as part of a single microDC. We plan to extend these models to deal with geo-distribution. The objective is to propose models that will enable us to refine our placement algorithms as discussed in the next paragraph. These models should be able

to consider the energy consumption induced by all WAN data exchanges, including site-to-site data movements as well as the end users' communications for accessing virtualized resources.

Third, we expect to implement green-energy-aware balancing strategies, leveraging the aforementioned contributions.

Although the infrastructures we envision increase complexity (because WAN aspects should also be taken into account), the geo-distribution of resources brings several opportunities from the energy viewpoint. For instance, it is possible to define several workload/data placement policies according to renewable energy availability. Moreover, a tightly-coupled software stack allows users to benefit from such a widely distributed infrastructure in a transparent way while enabling administrators to balance resources in order to benefit from green energy sources when available.

An important difficulty, compared to centralized infrastructures, is related to data sharing between software instances. In particular, we will study issues raised by the distribution and replication of services across several microDCs. In this new context, many challenges must be addressed: where to place the data (Cloud, Edge) in order to mitigate data movements? What is the impact in terms of energy consumption, network and response time of these two approaches? How to manage the consistency of replicated data/services? All these aspects must be studied and integrated into our placement algorithms.

Fourth, we will investigate the energy footprint of the current techniques that address failure and performance variability in large-scale systems. For instance, *stragglers* (*i.e.*, tasks that take a significantly longer time to finish than the normal execution time) are natural results of performance variability, they cause extra resource and energy consumption. Our goal is to understand the energy overhead of these techniques and introduce new handling techniques that take into consideration the energy efficiency of the platform [79].

Finally, in order to answer specific energy constraints, we want to reify energy aspects at the application level and propose a metric related to the use of energy (Green SLA [27]), for example to describe the maximum allowed CO₂ emissions of a Fog/Edge service. Unlike other approaches [60, 32, 58] that attempt to identify the best trade-off, we want to offer to developers/end-users the opportunity to select the best choice between application performance, correctness and energy footprint. Such a capability will require reifying the energy dimension at the level of big-data and interactive applications. Besides, with the emergence of renewable energy (*e.g.*, solar panels for microDC), investigating the energy consumption *vs* performance trade-off [63] and the smart usage of green energy for ICT geo-distributed services seems promising. For example, we want to offer the opportunity to developers/end-users to control the scaling of the applications based on this trade-off instead of current approaches that only considered application load. Providing such a capability will also require appropriate software abstractions.

3.5 Security

Because of its large size and complex software structure, geo-distributed applications and infrastructures are particularly exposed to security and privacy issues [83]. They are subject to numerous security vulnerabilities that are frequently exploited by malicious attackers in order to exfiltrate personal, institutional or corporate data. Securing these systems require security and privacy models and corresponding techniques that are applicable at all software layers in order to guard interactions at each level but also between levels. However, very few security models exist for the lower layers of the software stack and no model enables the handling of interactions involving the complete software stack. Any modification to its implementation, deployment status, configuration, etc., may introduce new or trigger existing security and privacy issues. Finally, applications that execute on top of the software stack may introduce security issues or be affected by vulnerabilities of the stack. Overall, security and privacy issues are therefore interdependent with all other activities of the STACK team and constitute an important research topic for the team.

As part of the STACK activities, we consider principally security and privacy issues related to the vertical and horizontal compositions of software components forming the software stack and the distributed applications running on top of it. Modifications to the *vertical composition* of the software stack affect different software levels at once. As an example, side-channel attacks often target virtualized services (*i.e.*, services running within VMs); attackers may exploit insecure hardware caches at the system level to exfiltrate data from computations at the higher level of VM services [77, 94]. Security and privacy issues also affect *horizontal compositions*, that is, compositions of software abstractions on one level: most

frequently horizontal compositions are considered on the level of applications/services but they are also relevant on the system level or the middleware level, such as compositions involving encryption and database fragmentation services.

The STACK members aim at addressing two main research issues: enabling full-stack (vertical) security and per-layer (horizontal) security. Both of these challenges are particularly hard in the context of large geo-distributed systems because they are often executed on heterogeneous infrastructures and are part of different administrative domains and governed by heterogeneous security and privacy policies. For these reasons they typically lack centralized control, are frequently subject to high latency and are prone to failures.

Concretely, we will consider two classes of security and privacy issues in this context. First, on a general level, we strive for a method for the programming and reasoning about compositions of security and privacy mechanisms including, but not limited to, encryption, database fragmentation and watermarking techniques. Currently, no such general method exists, compositions have only been devised for specific and limited cases, for example, compositions that support the commutation of specific encryption and watermarking techniques [69, 41]. We provided preliminary results on such compositions [42] and have extended them to biomedical, notably genetic, analyses in the e-health domain [34]. Second, on the level of security and privacy properties, we will focus on isolation properties that can be guaranteed through vertical and horizontal composition techniques. We have proposed first results in this context in form of a compositional notion of distributed side channel attacks that operate on the system and middleware levels [30].

It is noteworthy that the STACK members do not have to be experts on the individual security and privacy mechanisms, such as watermarking and database fragmentation. We are, however, well-versed in their main properties so that we can integrate them into our composition model. We also interact closely with experts in these techniques and the corresponding application domains, notably e-health for instance, in the context of the PRIVGEN project³, see Section 10.

More generally, we highlight that security issues in distributed systems are very closely related to the other STACK challenges, dimensions and research directions. Guaranteeing security properties across the software stack and throughout software layers in highly volatile and heterogeneous geo-distributed systems is expected to harness and contribute results to the self-management capabilities investigated as part of the team's resource management challenges. Furthermore, security and privacy properties are crosscutting concerns that are intimately related to the challenges of application life cycle management. Similarly, the security issues are also closely related to the team's work on programming support. This includes new means for programming, notably in terms of event and stream programming, but also the deployment and reconfiguration challenges, notably concerning automated deployment. As a crosscutting functionality, the security challenges introduced above must be met in an integrated fashion when designing, constructing, executing and adapting distributed applications as well as managing distributed resources.

4 Application domains

4.1 Overview

Supporting industrial actors and open-source communities in building an advanced software management stack is a key element to favor the advent of new kinds of information systems as well as web applications. Augmented reality, telemedicine and e-health services, smart-city, smart-factory, smart-transportation and remote security applications are under investigations. Although, STACK does not intend to address directly the development of such applications, understanding their requirements is critical to identify how the next generation of ICT infrastructures should evolve and what are the appropriate software abstractions for operators, developers and end-users. STACK team members have been exchanging since 2015 with a number of industrial groups (notably Orange Labs and Airbus), a few medical institutes (public and private ones) and several telecommunication operators in order to identify both opportunities and challenges in each of these domains, described hereafter.

³Privacy-preserving sharing and processing of genetic data

4.2 Industrial Internet

The Industrial Internet domain gathers applications related to the convergence between the physical and the virtual world. This convergence has been made possible by the development of small, lightweight and cheap sensors as well as complex industrial physical machines that can be connected to the Internet. It is expected to improve most processes of daily life and decision processes in all societal domains, affecting all corresponding actors, be they individuals and user groups, large companies, SMEs or public institutions. The corresponding applications cover: the improvement of business processes of companies and the management of institutions (*e.g.*, accounting, marketing, cloud manufacturing, etc.); the development of large “smart” applications handling large amounts of geo-distributed data and a large set of resources (video analytics, augmented reality, etc.); the advent of future medical prevention and treatment techniques thanks to the intensive use of ICT systems, etc. We expect our contributions will favor the rise of efficient, correct and sustainable massively geo-distributed infrastructures that are mandatory to design and develop such applications.

4.3 Internet of Skills

The Internet of Skills is an extension of the Industrial Internet to human activities. It can be seen as the ability to deliver physical experiences remotely (*i.e.*, via the Tactile Internet). Its main supporters advocate that it will revolutionize the way we teach, learn, and interact with pervasive resources. As most applications of the Internet of Skills are related to real time experiences, latency may be even more critical than for the Industrial Internet and raise the locality of computations and resources as a priority. In addition to identifying how Utility Computing infrastructures can cope with this requirement, it is important to determine how the quality of service of such applications should be defined and how latency and bandwidth constraints can be guaranteed at the infrastructure level.

4.4 e-Health

The e-Health domain constitutes an important societal application domain of the two previous areas. The STACK teams is investigating distribution, security and privacy issues in the fields of systems and personalized (aka. precision) medicine. The overall goal in these fields is the development of medication and treatment methods that are tailored towards small groups or even individual patients.

We are working, as part of the ongoing PrivGen CominLabs collaborative project on new means for the sharing of genetic data and applications in the Cloud. More generally, we are applying and developing corresponding techniques for the medical domains of genomics, immunobiology and transplantology in the international network SHLARC and the regional networks SysMics and Oncoshare (see Section 10): there, we investigate how to secure and preserve privacy if potentially sensitive personal data is moved and processed by distributed biomedical analyses.

We are also involved in the SyMeTRIC regional initiative where preliminary studies have been conducted in order to build a common System Medicine computing infrastructure to accelerate the discovery and validation of bio-markers in the fields of oncology, transplantation, and chronic cardiovascular diseases. The challenges were related to the need of being able to perform analyses on data that cannot be moved between distinct locations.

The STACK team will continue to contribute to the e-Health domain by harnessing advanced architectures, applications and infrastructures for the Fog/Edge.

4.5 Network Virtualization and Mobile Edge Services

Telecom operators have been among the first to advocate the deployment of massively geo-distributed infrastructures, in particular through working groups such as the [Mobile Edge Computing at the European Telecommunication Standards Institute](#). The initial reason is that geo-distributed infrastructures will enable Telecom operators to virtualize a large part of their resources and thus reduce capital and operational costs. As an example, we are investigating through the I/O Lab, the joint lab between Orange and Inria, how can a Centralized Radio Access Networks (*a.k.a.* C-RAN or Cloud-RAN) be supported for 5G networks. We highlight that our expertise is not on the network side but rather on where and how we

can deploy, allocate and reconfigure software components, which are mandatory to operate a C-RAN infrastructure, in order to guarantee the quality of service expected by the end-users. Finally, working with actors from the network community is a valuable advantage for a distributed system research group such as STACK. Indeed, achievements made within one of the two communities serve the other.

5 Social and environmental responsibility

5.1 Footprint of research activities

In addition to the international travels⁴, the environmental footprint of our research activities is linked to our intensive use of large-scale testbeds such as Grid'5000 (STACK members are often in the top10 list of the largest consumers). Although the access to such facilities is critical to move forward in our research roadmap, it is important to recognize that they have a strong environmental impact as described in the next paragraph.

5.2 Impact of research results

The environmental impact of digital technology is a major scientific and societal challenge. Even though the software remains virtual objects, it is executed on very real hardware contributing to the carbon footprint. This impact materializes during the manufacture / destruction of hardware infrastructure (estimated at 45% of digital consumption in 2018 by the The Shift Project) and during the software use phase via terminals, networks and data centers (estimated at 55%). Stack members have been studying various approaches for several years to reduce the energy footprint of digital infrastructures during the use phase. The work carried out revolves around two main axes: (i) reducing the energy footprint of infrastructures and (ii) adapting the software applications hosted by these infrastructures according to the energy available. More precisely, this second axe investigates possible improvements that could be made by the end-users of the software themselves. At scale, involving end-users in decision-making processes concerning energy consumption would lead to more frugal Cloud computing. For instance, in the GL4MA project (cf. Section 10), we propose that the end-users customize the software in SaaS mode to contribute to reducing their carbon footprint by using either fewer resources or resources powered directly by renewable energy.

6 Highlights of the year

Regarding scientific results, the team has produced a number of outstanding results on the management of resources and data in large-scale infrastructures. In particular, the team published a survey on decentralized control planes for Fog/Edge infrastructures [3]. Furthermore the journal paper published at Science of Computer Programming brings an important stone to the results of the team on distributed software deployment and reconfiguration [1]. We should also mention the extension of our activities to the IoT area with the recent recruiting of Remous-Aris Koutsiamanis, Asso. Prof. who joined the team in Oct 2020. Among others, we have concluded work studying and achieving energy-efficient and reliable low-power wireless IoT communications [4].

On the software side, the team has pursued its efforts on the development of the **EnosLib library** and the resulting artifacts to help researchers perform experiment campaigns. In particular, we would like to point out the recent extensions toward the Kubernetes framework as well as the recent article we published in 2021 [2].

Finally, on the platform side, we continued our effort and took part to the different actions around the SILECS initiative (and its European part, SLICES), see Section 10.

6.1 Awards

In 2021, the team has received the Best student paper award at IEEE International Conference on Computers and Communications (IEEE ISCC 2021) for its paper: « The cost of immortality: a time to live

⁴Due to the COVID pandemia, STACK members did not travel this year overall.

for smart contracts. » by Dimitri Saingre, Thomas Ledoux et Jean Marc Menaud.

We would like also to highlight additional elements that underline the visibility and recognition of the team nationally and internationally. First, at the national level, the team has obtained an important grant from the PIA 4 program "Appel à manifestation d'intérêt relatif à la Stratégie d'accélération Cloud" for the OTPaaS project (56M€ with 1.2M€ for the team). The OTPaaS project targets the design and development of a complete software stack to administrate and use edge infrastructures for the industry sector. The consortium brings together national and user technology suppliers from major groups (Atos / Bull, Schneider Electric, Valeo) and SMEs / ETIs (Agileo Automation, Mydatamodels, Dupliprint, Solem, Tridimeo, Prosyst, Soben), with a strong support from major French research institutes (CEA, Inria, IMT, CAPTRONIC).

At the international level, the involvement of the team in the OpenStack community has been officially recognized: IMT Atlantique has been invited to officially become one of the associate members of the Open Infrastructure foundation (see [the OpenInfra website](#)).

7 New software and platforms

STACK's members are initiators and contributors of multiple software development as well as platforms. We present in this section the major ones.

7.1 New software

7.1.1 MAD

Name: Madeus Application Deployer

Keywords: Automatic deployment, Distributed Software, Component models, Cloud computing

Scientific Description: MAD is a Python implementation of the Madeus deployment model for multi-component distributed software. Precisely, it allows to: 1. describe the deployment process and the dependencies of distributed software components in accordance with the Madeus model, 2. describe an assembly of components, resulting in a functional distributed software, 3. automatically deploy the component assembly of distributed software following the operational semantics of Madeus.

Functional Description: MAD is a Python implementation of the Madeus deployment model for multi-component distributed software. Precisely, it allows to: 1. describe the deployment process and the dependencies of distributed software components in accordance with the Madeus model, 2. describe an assembly of components, resulting in a functional distributed software, 3. automatically deploy the component assembly of distributed software following the operational semantics of Madeus.

Release Contributions: Initial submission with basic functionalities of MAD

News of the Year: Operational prototype.

Publications: [hal-01858150](#), [hal-01897803](#)

Contact: Hélène Coullon

Participants: Christian Perez, Dimitri Pertin, Hélène Coullon, Maverick Chardet

Partners: IMT Atlantique, LS2N, LIP

7.1.2 Nitro

Keywords: Cloud storage, Virtual Machine Image, Geo-distribution

Scientific Description: Nitro is a storage system that is designed to work in geo-distributed cloud environments (i.e., over WAN) to efficiently manage Virtual Machine Images (VMIs).

Nitro employs fixed-size deduplication to store VMIs. This technique contributes to minimizing the network cost. Also, Nitro incorporates a network-aware scheduling algorithm (based on max flow algorithm) to determine which chunks should be pulled from which site in order to reconstruct the corresponding image on the destination site, with minimal (provisioning) time.

Functional Description: Geo-distributed Storage System to optimize Images (VM, containers, ...) management, in terms of cost and time, in geographically distributed cloud environment (i.e. data centers are connected over WAN).

URL: <https://gitlab.inria.fr/jdarrous/nitro>

Authors: Jad Darrous, Shadi Ibrahim, Christian Perez

Contact: Shadi Ibrahim

7.1.3 VMPlaceS

Keywords: Simulation, Virtualization, Scheduling

Functional Description: VMPlaces is a dedicated framework to evaluate and compare VM placement algorithms. This framework is composed of two major components: the injector and the VM placement algorithm. The injector is the generic part of the framework (i.e. the one you can directly use) while the VM placement algorithm is the part you want to study (or compare with available algorithms). Currently, the VMPlaceS is released with three algorithms:

Entropy, a centralized approach using a constraint programming approach to solve the placement/reconfiguration VM problem

Snooze, a hierarchical approach where each manager of a group invokes Entropy to solve the placement/reconfiguration VM problem. Note that in the original implementation of Snooze, it is using a specific heuristic to solve the placement/reconfiguration VM problem. As the sake of simplicity, we have simply reused the entropy scheduling code.

DVMS, a distributed approach that dynamically partitions the system and invokes Entropy on each partition.

URL: <http://beyondtheclouds.github.io/VMPlaceS/>

Contact: Adrien Lebre

Participants: Adrien Lebre, Jonathan Pastor, Mario Südholt

7.1.4 ENOS

Name: Experimental eNvironment for OpenStack

Keywords: OpenStack, Experimentation, Reproducibility

Functional Description: Enos workflow :

A typical experiment using Enos is the sequence of several phases: - enos up : Enos will read the configuration file, get machines from the resource provider and will prepare the next phase - enos os : Enos will deploy OpenStack on the machines. This phase rely highly on Kolla deployment. - enos init-os : Enos will bootstrap the OpenStack installation (default quotas, security rules, ...) - enos bench : Enos will run a list of benchmarks. Enos support Rally and Shaker benchmarks. - enos backup : Enos will backup metrics gathered, logs and configuration files from the experiment.

URL: <https://github.com/BeyondTheClouds/enos>

Contact: Adrien Lebre

Partner: Orange Labs

7.1.5 EnOSlib

Name: EnOSlib is a library to help you with your experiments

Keywords: Distributed Applications, Distributed systems, Evaluation, Grid Computing, Cloud computing, Experimentation, Reproducibility, Linux, Virtualization

Functional Description: EnOSlib is a library to help you with your distributed application experiments. The main parts of your experiment logic is made reusable by the following EnOSlib building blocks:

- Reusable infrastructure configuration: The provider abstraction allows you to run your experiment on different environments (locally with Vagrant, Grid'5000, Chameleon and more)
- Reusable software provisioning: In order to configure your nodes, EnOSlib exposes different APIs with different level of expressivity
- Reusable experiment facilities: Tasks help you to organize your experimentation workflow.

EnOSlib is designed for experimentation purpose: benchmark in a controlled environment, academic validation ...

URL: <https://discovery.gitlabpages.inria.fr/enoslib/>

Publications: [hal-01664515](#), [hal-01689726](#)

Contact: Mathieu Simonin

7.1.6 Concerto

Name: Concerto

Keywords: Reconfiguration, Distributed Software, Component models, Dynamic software architecture

Scientific Description: Concerto is a reconfiguration model which allows to describe distributed software as an evolving assembly of components.

Functional Description: Concerto is an implementation of the formal model Concerto written in Python. Concerto allows to : 1. describe the life-cycle and the dependencies of software components, 2. describe a components assembly that forms the overall life-cycle of a distributed software, 3. automatically reconfigure a Concerto assembly of components by using a set of reconfiguration instructions as well as a formal operational semantics.

News of the Year: In 2020, we added the ability to read from and write to non-data provide ports. We also updated the Madeus wrapper to the new version with only USE and PROVIDE ports and no groups.

URL: <https://gitlab.inria.fr/VerDi-project/concerto>

Publications: [hal-03103714](#), [hal-02535077](#), [hal-01897803](#)

Contact: H el ene Coullon

Partners: IMT Atlantique, LS2N, LIP

7.2 New platforms

7.2.1 OpenStack

OpenStack is the de facto open-source management system to operate and use Cloud Computing infrastructures. Started in 2012, the OpenStack foundation gathers 500 organizations including groups such as Intel, AT&T, RedHat, etc. The software platform relies on tens of services with a 6-month development cycle. It is composed of more than 2 millions of lines of code, mainly in Python, just for the core services. While these aspects make the whole ecosystem quite swift, they are also good signs of maturity of this community.

We created and animated between 2016 and 2018 the **Fog/Edge/Massively Distributed (FEMDC) Special Interest Group** and have been contributing to the Performance working group since 2015. The former investigates how OpenStack can address Fog/Edge Computing use cases whereas the latter addresses scalability, reactivity and high-availability challenges. In addition to releasing white papers and guidelines [89], the major result from the academic view point is the aforementioned EnOS solution, a holistic framework to conduct performance evaluations of OpenStack (control and data plane). In May 2018, the FEMDC SiG turned into a larger group under the control of the OpenStack foundation. This group gathers large companies such as Verizon, ATT, etc. Although our involvement has been less important in 2020, our participation is still significant. For instance, we co-signed the second white paper delivered by the edge working group in 2020 [90].

7.2.2 Grid'5000

Grid'5000 is a large-scale and versatile testbed for experiment-driven research in all areas of computer science, with a focus on parallel and distributed computing including Cloud, HPC and Big Data. It provides access to a large amount of resources: 12000 cores, 800 compute-nodes grouped in homogeneous clusters, and featuring various technologies (GPU, SSD, NVMe, 10G and 25G Ethernet, Infiniband, Omni-Path) and advanced monitoring and measurement features for traces collection of networking and power consumption, providing a deep understanding of experiments. It is highly reconfigurable and controllable. STACK members are strongly involved into the management and the supervision of the testbed, notably through the steering committee or the SeDuCe testbed described hereafter.

7.2.3 SeDuCe

The SeDuCe Project aims to deliver a research testbed dedicated to holistic research studies on energetical aspects of datacenters. Part of the Grid'5000 Nantes' site, this infrastructure is composed of probes that measure the power consumption of each server, each switch and each cooling system, and also measure the temperature at the front and the back of each servers. These sensors enable reasearch to cover a full spectrum of the energetical aspect of datacenters, such as cooling and power consumption depending of experimental conditions.

The testbed is connected to renewable energy sources (solar panels). This "green" datacenter enables researchers to perform real experiment-driven studies on fields such as temperature based scheduling or "green" aware software (*i.e.*, software that take into account renewable energies and weather conditions).

7.2.4 PiSeDuCe

In 2021, we consolidated the development of PiSeDuCe, a deployment and reservation system for Edge Computing infrastructures composed of multiple Raspberry Pi Cluster started in 2020. Typically, a cluster of 8 Raspberry Pi costs less than 900 euros and only needs an electrical outlet and a wifi connection for its installation and configuration. Funded by the CNRS through the Kabuto project, and in connection with the SILECS initiative, we have extended PiSeducer to propose a device to cloud deployment system (from devices on Fit IoT Lab to servers in Grid'5000).

7.2.5 SILECS/SLICES

STACK Members are involved in the definition and bootstrap of the SILECS infrastructure. This infrastructure can be seen as a merge of the Grid'5000 and FIT testbeds with the goal of providing a common

platform for experimental computer Science (Next Generation Internet, Internet of things, clouds, HPC, big data, etc.). In 2021, STACK contributions are mainly related to the European initiatives relative to SILECS. In particular, members have taken part to the SLICES-DS project (Design Study) as well as the submission of SLICES-PP (Preparatory Phase) of the [SLICES-RI action](#).

8 New results

8.1 Resource Management

Participants: Ronan-Alexandre Cherrueau, Marie Delavergne, David Espinel, Remous Aris Koutsiamanis, Adrien Lebre, Brice Nedelec.

The evolution of the cloud computing paradigm in the last decade has amplified the access of on-demand services (economically attractive, easy-to-use manner, etc.). However, the current model built upon a few large datacenters (DCs) may not be suited to guarantee the needs of new use cases, notably the boom of the Internet of Things (IoT). To better respond to the new requirements (in terms of delay, traffic, etc.), compute and storage resources should be deployed closer to the end-user. In the case of telecommunication operators, the network Point of Presence (PoP), which they have always operated, can be inexpensively extended to host these resources. The question is then how to manage such a massively Distributed Cloud Infrastructure (DCI) to provide end-users the same services that made the current cloud computing model so successful. In 2021, we have continued our effort to answer this question and delivered multiple contributions.

Resource management systems/middleware for Edge infrastructures: The two first contributions have been conducted within the framework of the Inria/Orange Laboratory. More precisely, we have investigated how connectivity among several resource managers in charge of operating, each one a subset of the infrastructure, could be established. In [3], we have surveyed and analyzed the characteristics and limitations of existing technologies in the Software Defined Network field that could be used to provide the intersite connectivity feature. We have also introduced how the network is addressed in the Kubernetes ecosystem, and have analyzed its use in the proposed context. We have concluded by providing a discussion about some research directions in the field of SDN applied to distributed Cloud-Edge infrastructures' management. Leveraging this survey, we propose the DIMINET solution [22], a service in charge of providing on-demand connectivity for multiple sites. DIMINET leverages a logically and physically distributed architecture where instances collaborate on-demand and with minimal traffic exchange to provide inter-site connectivity management. The lessons learned during this study allows us to propose the premises of a generalization in order to be able to distribute in a non-intrusive manner any service in a DCI.

In [16], we have extended the aforementioned study on the vanilla Kubernetes framework with in-vivo experiments. Cloud Computing has highlighted the importance of container orchestration such as Kubernetes to manage the life-cycle of distributed applications. With the advent of Edge Computing, DevOps expect to find the features of containers in the cloud, also at the edge. However, similarly to systems such as OpenStack, orchestration solutions have not been designed to deal with geo-distribution aspects such as latency, intermittent networks, etc. In other words, it is unclear whether they could be directly used on top of massively distributed edge infrastructures without revision. To answer this question, we have conducted an evaluation of Kubernetes in a WANWide context leveraging the Grid'5000 testbed. More precisely, we have discussed results we obtained during an experimental campaign that aimed at analyzing the impact of WAN links on its behaviour. While multiple initiatives investigate how Kubernetes could be revised to deal with distribution aspects, this work, to the best of our knowledge, is the first one to rigorously evaluate whether the Kubernetes vanilla code could be directly used without any change. In [24], we investigated how peers in a geo-distributed infrastructure could maintain an index of relevant replicas. Although the holy grail of storing and manipulating data in Edge infrastructures is yet to be found, state-of-the-art approaches demonstrated the relevance of replication strategies

that bring content closer to consumers: The latter enjoy better response time while the volume of data passing through the network decreases overall. Unfortunately, locating the closest replica of a specific content requires indexing every live replica along with its location. Relying on remote services enters in contradiction with the properties of Edge infrastructures as locating replicas may effectively take more time than actually downloading content. At the opposite end, maintaining such an index at every node would prove overly costly in terms of memory and traffic, especially since nodes can create and destroy replicas at any time. We proposed to abstract this content indexing challenge as a distributed partitioning problem: every node only indexes its closest replica, and connected nodes with a similar index compose a partition. Our decentralized implementation AS-cast is (i) efficient, for it uses partitions to lock down the traffic generated by its operations to relevant nodes, yet it (ii) guarantees that every node eventually acknowledges its partition despite concurrent operations. Our complexity analysis supported by simulations shows that AS-cast scales well in terms of generated traffic and termination time. As such, AS-cast can constitute a new building block for geo-distributed services.

IoT resource management: As aforementioned, the team has initiated new activities covering the Cloud-IoT continuum.

In [10], we investigated the challenges and limitations in using a distributed network resource allocation mechanism in the wireless Industrial IoT context. The wireless communication protocol IEEE Std 802.15.4-2015 Time Slotted Channel Hopping (TSCH) is the de facto Medium Access Control (MAC) mechanism for industrial applications. It renders communications more resilient to interference by spreading them over the time (time-slotted) and the frequency (channel-hopping) domains. The 6TiSCH architecture bases itself on this new MAC layer to enable high reliability communication in Wireless Sensor Networks (WSNs). In particular, it manages the construction of a distributed communication schedule that continuously adapts to changes in the network. Our investigation comprises first a thorough description of the 6TiSCH architecture, the 6TiSCH Operation Sublayer (6top), and the Minimal Scheduling Function (MSF). We then study its behavior and reactivity from low to high traffic rates by employing the Python-based 6TiSCH simulator. Our performance evaluation results demonstrate that the convergence pattern of MSF is the root cause of the majority of packet losses observed in the network. We also show that MSF is prone to over-provisioning of the network resources, especially in the case of varying traffic load. We propose a mathematical model to predict the convergence pattern of MSF. Finally we investigate the impact of varying parameters on the behavior of the scheduling function.

In [6], we investigate the trade-offs of computational, memory and network traffic resource consumption against reliability of different packet fragmentation and forwarding mechanisms in the wireless Industrial IoT context which places significant demands in terms of reliability on wireless connectivity. The IEEE Std 802.15.4-2015 standard was designed in response to these demands, and the IPv6 over Low power Wireless Personal Area Networks (6LoWPAN) adaptation layer was introduced to address (among other issues) its payload size limitations by performing packet compression and fragmentation. However, the standardised method does not cope well with low link-quality situations and, thus, we present the state-of-the-art Forward Error Correction (FEC) methods and introduce our own contribution, Network Coding FEC (NCFEC), to improve performance in these situations. We present and analyse the existing methods as well as our own theoretically, and we then implement them and perform an experimental evaluation using the 6TiSCH simulator. The simulation results demonstrate that when high reliability is required and only low quality links are available, NCFEC performs best, with a trade-off between additional network and computational overhead. In situations where the link quality can be guaranteed to be higher, simpler solutions also start to be feasible, but with reduced adaptation flexibility.

Experiment-driven studies: The last contribution we achieved is related to our experiment-driven studies we have performed since 2016. Despite the importance of such activities in the distributed computing community, there has been little progress in helping researchers conduct their experiments. In most cases, we have to achieve tedious and time-consuming development and instrumentation activities to deal with the specifics of testbeds and the system under study. In order to relieve researchers of the burden of those efforts, we have developed ENOSLIB: a Python library that takes into account best experimentation practices and leverages modern toolkits on automatic deployment and configuration systems [2]. ENOSLIB helps researchers not only in the process of developing their experimental artifacts,

but also in running them over different infrastructures. To demonstrate the relevance of our library, we discuss three experimental engines built on top of ENOSLIB, and used to conduct empirical studies on complex software stacks between 2016 and 2019 (database systems, communication buses and Open-Stack). By introducing ENOSLIB, our goal is to gather academic and industrial actors of our community around a library that aggregates everyday experiment-driven research operations. A library that has been already adopted by open-source projects and members of the scientific community thanks to its ease of use and extension (see Section 7 for further details).

8.2 Programming Support

Participants: Ronan-Alexandre Cherrueau, H el ene Coullon, Marie Delavergne, Adrien Lebre, Jolan Philippe, Simon Robillard, Emile Cadorel.

Model-Driven Engineering (MDE): MDE is a software programming approach that raises the level of abstraction in traditional programming languages by using models with recurring design patterns. MDE simplifies the collaboration between development teams and more broadly promotes compatibility between systems. MDE is increasingly used to help in the development of distributed systems, microservices architectures, and IoT systems but is also leveraged, for instance, in lowcode platforms such as Node-RED to write IoT applications. In [17] two important drawbacks of MDE have been studied. First, as most of the high-abstraction level programming models, when dealing with very big models most MDE frameworks become poorly efficient and slow, thus loosing their initial advantage of speeding up the development process. Second, because of the underlying complexity of MDE approaches the level of confidence in MDE frameworks is usually low. In this paper, we extend the formal transformation engine (transformation being an important part of MDE approaches) CoqTL with a more scalable implementation that is proven equivalent to the initial one, and that is automatically dumped to a Apache Spark code. We have developed a prototype implementation of the refined specification on top of Spark, and have evaluated its performance on a simple case study to assess the speedup our solution can reach.

Software deployment/reconfiguration: If speeding and simplifying the development process of distributed systems is of prior importance because of the increasing complexity of geographically distributed infrastructures, speeding and simplifying the deployment of these systems, and enabling and speeding their dynamic evolution through time is another key to making these infrastructures usable. Indeed, because of the scale of both geo-distributed infrastructures and distributed systems (e.g., microservices architectures) manually handling deployments and reconfiguration procedures is an error-prone tedious and complex task, probably impossible in practice. Hence, these procedures need to be automated, or better autonomous. Reconfiguration of systems is a difficult and sensitive procedure during which the system is in between two states, the previous state and the desired state. For this reason, we are interested in two metrics in particular for a few years: the efficiency of the reconfiguration to reach as quick as possible the desired new state; and the safety of the reconfiguration to obtain some formal guarantees on the procedure. In the journal paper [1] we extensively present the reconfiguration formal model Concerto which poses an important pillar in our journey towards autonomous, efficient and safe reconfiguration. Concerto is used to manage the lifecycle of software components and coordinate their reconfiguration operations. Concerto promotes efficiency with a fine-grained representation of dependencies and parallel execution of reconfiguration actions, both within components and between them. In this paper, the elements of the model are described as well as their formal semantics. In addition, we outline a performance model that can be used to estimate the time required by reconfigurations, and we describe an implementation of the model. The evaluation demonstrates the accuracy of the performance estimations, and illustrates the performance gains provided by the execution model of Concerto compared to state-of-the-art systems.

Reifying geo-distribution at the software level: Following the results obtained during David Espinel's PhD (see Section 8.1 [22]), we have investigated how a software stack that has been designed to run

on one DC, could benefit from geo-distribution (locality) while dealing with the inherent constraints of wide-area network links without requiring intrusive changes? The admitted approach for several years consists in modifying cloud applications by entangling geo-distribution aspects in the business logic using distributed data stores. However, this makes the code intricate and contradicts the software engineering principle of externalizing concerns. In [15], and then [20], we propose a different approach that relies on the modularity property of microservice applications: (i) one instance of an application is deployed at each edge location, making the system more robust to network partitions (local requests can still be satisfied), and (ii) collaboration between instances can be programmed outside of the application in a generic manner thanks to a service mesh. We validate the relevance of our proposal on a real use-case: geo-distributing OpenStack, a modular application composed of 13 million of lines of code and more than 150 services. We underline these results are at the frontier of the two main research activities of the STACK team (*i.e.*, resource management and programming support).

8.3 Energy-aware computing

Participants: Remous Aris Koutsiamanis, Thomas Ledoux, Jean-Marc Menaud, Dimitri Saingre.

Blockchain technologies: For the past few years, we have been investigating the energy consumption of blockchains, which is particularly criticized. Unfortunately, energy evaluation is often performed with ad hoc tools and specific experimental environments. We have therefore developed BCTMark, a generic framework for benchmarking blockchain technologies on an emulated network in a reproducible way. Leveraging the aforementioned EnosLib library (cf. Section 7.1), BCTMark has enabled us to study a key aspect of the energy consumption of modern blockchains: smart-contract execution [12].

Smart contracts, scripts at the heart of blockchain-based applications, are meant to be available forever once deployed. However, this property comes at a price: the space required to store new contracts continues to increase. We show that over the course of a year, 70% of deployed contracts are not used while they continue to occupy space on the blockchain. To tackle this issue, we have proposed a new protocol to identify and delete unused contracts. Through simulation, based on historical Ethereum data, we have shown that deleting smart contracts after a 90-day period of inactivity could lead to a 66% reduction in the number of contracts stored over a year [5].

Energy efficiency of Industrial IoT devices: In [4], we focus on the trade-off between energy-consumption, reliability and latency of different packet forwarding protocols in the wireless Industrial IoT context and propose a new a "braided" forwarding pattern. The aim is to improve existing low-power wireless communication protocols by supporting the strict Quality of Service (QoS) requirements of the industry. Focusing on the IEEE Std 802.15.4-2015 TSCH link-layer standard and the RPL standard at the IETF, we present On-Demand Selection (ODeSe), a novel multi-path routing algorithm, which improves our previous work, the Common Ancestor (CA) algorithms, by selecting the most suitable upward forwarders. Using the Cooja network simulator running the Contiki OS, we compare ODeSe against single-path RPL and multi-path RPL with different alternative parent selection algorithms. The results demonstrate that ODeSe outperforms single-path RPL in terms of reliability, and multi-path RPL in terms of energy consumption while maintaining a 99.14% packet delivery ratio.

8.4 Security and Privacy

Participants: Maxime Belair, Fatima-zahra Boujdad, Wilmer Edicson Garzon Alfonso, Jean Marc Menaud, Sirine Sayadi, Mario Südholt.

Secure and privacy-preserving distributed genetic analyses: The increasing availability of sequenced human genomes is enabling health professionals and genomics researchers to well understand the implication of genetic variants in the development of common diseases, notably by means of genome-wide association studies (GWAS) which are very promising for personalized medicine and diagnostic testing. However, the need to handle genetic data from different sources in order to conduct large studies entails multiple privacy and security issues. Actually, classical methods of anonymization are inapplicable for genetic data that are now known to be identifying per se.

We have proposed a novel framework for privacy-preserving collaborative GWAS performed in the cloud. Indeed, our proposal is the first framework which combines a hybrid cloud deployment with a set of four security mechanisms: digital watermarking, homomorphic encryption, meta-data de-identification and the Intel Software Guard Extensions technology. Using these mechanism we can ensure confidentiality of genetic data as well as their integrity. Furthermore, our approach describes meta-data management which has rarely been considered in state-of-the-art propositions despite their importance to genetic analyses. In addition, the new deployment model we have suggested fits with existing infrastructures which makes its integration straightforward. Experimental results of a prototypical implementation on typical data sizes have demonstrated that our solution protocol is feasible and that the framework is practical for real-world scenarios. [14]

Privacy-aware distributed FAM: FAMD analyses are an important statistical technique that not only enables the visualization of large data but also helps to select subgroups of relevant information for a given patient. While such analyses are well-known in the medical domain, they have to satisfy new data governance constraints if reference data is distributed, notably in the context of large consortia developing the future generation of analyses in the domain of personalised medicine.

We have motivated the use of distributed implementations for FAMD analyses in the context of the development of a personalised medicine application called KITAPP. We have presented a new distribution method for FAMD and evaluated its implementation in a multi-site setting based on real data. Finally we have studied how individual reference data is used to substantiate decision making, while enforcing a high level of usage control and data privacy for patients. [19]

Security properties of container technologies: Compared to full virtualization, containerization reduces virtualization overhead and resource usage, offers reduced deployment latency and improves reusability. For these reasons, containerization is massively used in an increasing number of applications. However, because containers share a full kernel with the host, they are more vulnerable to attacks that may compromise the host and the other containers on the system. In [13], we present SNAPPY (Safe Namespaceable And Programmable PolicY), a new framework that allows even unprivileged processes such as containers to safely and dynamically enforce in the kernel fine-grained, stackable and programmable eBPF security policies at runtime. This is done by making working coordinately a new LSM (Linux Security Module) Module, a new security Linux namespace abstraction (policy_NS) and eBPF policies enriched with 'dynamic helpers'. This design especially allows to minimize containers' attack surface.

9 Bilateral contracts and grants with industry

9.1 Bilateral contracts with industry

Participants: Ronan-Alexandre Cherrueau , Marie Delavergne , Adrien Lebre (*Contact point*) .

Following the ENOS bilateral contract (“Contrat de Recherche Externalisé”) between Orange and Inria (Sep 2017-Oct 2018), we agreed with Orange Labs to pursue this collaboration around a second contrat. This contrat (18 months for a budget of 150K€ ended mid of 2021. The main results are:

- A consolidation and extension of the Enos framework and the resulting EnosLib solution (see Section 6.4 and Section 6.5).
- An evaluation of Kubernetes in a WANWide context (which have led to a new experimental engine to evaluate multiple dimensions of the kubernetes ecosystem (for further information, please refer to [enos-kubernetes](#)).
- A new approach to reify location aspects at the CLI level in order to create new resources (image, VM, etc.) through a set of OpenStack instances while guaranteeing a notion of master copy (see Section 8).

9.2 Bilateral grants with industry

Alterway

Participants: Thomas Ledoux , Hiba Awad .

In 2020, during the preparation of the ANR SeMaFoR project, we started a cooperation with Alterway (www.alterway.fr), an SME specialized in Cloud and DevOps technologies. This cooperation resulted in a joint PhD thesis (called Cifre) entitled "A model-based approach for dynamic, multi-scale distributed systems" started in Nov. 2021.

OVHcloud

Participants: Thomas Ledoux , Pierre Jacquet .

In 2021, INRIA and OVHcloud have signed an agreement to jointly study the problem of a more frugal Cloud. They have identified 3 axes : (i) Software eco-design of Cloud services and applications; (ii) Energy efficiency leverages; (iii) Impact reduction and support for Cloud users. The Stack team obtained a PhD grant and a 24-month post-doc grant. The PhD student started last October, under a co-supervision with the Spirals team with the subject "Fostering the Frugal Design of Cloud Native Applications".

10 Partnerships and cooperations

10.1 International research visitors

Adjunct professorship in Norway

Participants: H el ene Coullon.

H el ene Coullon has been hired for two years from September 2020 to September 2022 as an adjunct professor at the Arctic University of Norway in Troms o. H el ene has given three keynotes to Master 2 students on distributed systems deployment and reconfiguration and has co-advised a Master 2 internship. She also collaborates with the local researchers Otto Anshus and Issam Rais with whom a joined PhD has begun on December 2021 (Antoine Omond).

10.2 European initiatives

10.2.1 FP7 & H2020 projects

Lowcomote

Participants: Hélène Coullon, Jolan Philippe.

Hélène Coullon is a member of the advisory board of the Lowcomote ITN project (H2020), on the subject of low code platforms. In particular, Hélène brings her expertise in distributed and high performance computing applied to model-driven engineering. She supervises Jolan Philippe, one PhD student of the project and a member of the team.

SLICES-DS

Participants: Adrien Lebre.

Adrien Lebre is a member of the SLICES Design Study project (ESFRI) that targets a Europe-wide test-platform designed to support large-scale, experimental research. It will provide advanced compute, storage and network components, interconnected by dedicated high-speed links. Pushing forward, the project's main goal is to strengthen the research excellence and innovation capacity of European researchers and scientists in the design and operation of digital infrastructures.

10.2.2 Other european programs/initiatives

SLICES-RI (ESFRI program)

Participants: Adrien Lebre.

Adrien Lebre is the IMT representative of the SLICES initiative (ESFRI) that targets a Europe-wide test-platform designed to support large-scale, experimental research on the cloud to IoT continuum.

10.3 National initiatives

10.3.1 Ademe

GLAMA

Participants: Brice Nédelec, Thomas Ledoux (*coordinator*).

The Green Label for Microservices Architecture (GLAMA) project aims to design and develop a technological platform (tools, framework, dedicated languages) for the self management of eco-responsible micro-service architectures for the Cloud. The experiments will be carried out through case studies provided by Sigma Informatique and the presence of renewable energy will initially be simulated. At the end of the project, the technological platform will be deployed as part of the CPER SeDuCe platform.

This project was founded by the Ademe (call Perfecto) running for 18 months (starting in September 2019) with an allocated budget of 116 480€ (the majority of aid was dedicated to the R&D engineer salary).

The work carried out during this project paves the way for new architectures and eco-designed applications involving Cloud service providers and users, in a virtuous way, around environmental issues. The final outcome of the project is still being evaluated. The increase in competence on this topic has led to initiate a PhD thesis in the context of the Inria-OVHcloud challenge (with the project-team Spirals) and a postdoc to come in the OTPaaS project.

10.3.2 CNRS

Kabuto

Participants: Remy Pottier, Jean-Marc Menaud (*coordinator*).

The Kabuto project aims to develop a software solution allowing the reservation and deployment IoT and Edge devices to HPC nodes. Strated in November 2020 for one year, Kabuto granted a post doc over one year (70 K€).

10.3.3 ANR

GRECO (ANR)

Participants: Adrien Lebre (*Contact point*), Matthieu Juzdzewski, Arnaud Szymanek.

The GRECO project (Resource manager for cloud of Things) was an ANR project (ANR-16-CE25-0016) running for 48 months (starting in January 2017 with an allocated budget of 522K€, 90K€ for STACK).

The consortium was composed of 4 partners: Qarnot Computing (coordinator) and 3 academic research group (DATAMOVE and AMA from the LIG in Grenoble and STACK from Inria Rennes Bretagne Atlantique).

The goal of the **GRECO project** is to design a manager for cloud of things. The manager should act at the IaaS, PaaS and SaaS layer of the cloud. To move forward to this objective, we have been designing a simulator to innovate in designing scheduling and data management systems. This simulator leverages the Simgrid/PyBATSIM solution.

SeMaFoR (Self-Management of Fog Resources)

Participants: Thomas Ledoux (*coordinator*), H el ene Coullon, Abdelghani Alidra.

Fog Computing is a paradigm that aims to decentralize the Cloud at the edge of the network to geographically distribute computing/storage resources and their associated services. It reduces bottlenecks and data movement. But managing a Fog is a major challenge: the system is larger, unreliable, highly dynamic and does not offer a global view for decision making. The objective of the SeMaFoR project is to model, design and develop a generic and decentralized solution for the self-management of Fog resources.

The consortium is composed of three partners: LS2N-IMT Atlantique (Stack, NaoMod, TASC), LIP6-Sorbonne Universit e (Delys), Alter way (SME). The Stack team supervises the project.

SeMaFoR is running for 42 months (starting in March 2021 with an allocated budget of 506k€, 230K€ for STACK).

PicNic (Transfert de grands volumes de donn ees entre datacenters)

Participants: Adrien Lebre, Jean Marc Menaud (*STACK representative*).

Large dataset transfer from one datacenter to another is still an open issue. Currently, the most efficient solution is the exchange of a hard drive with an express carrier, as proposed by Amazon with its SnowBall offer. Recent evolutions regarding datacenter interconnects announce bandwidths from 100 to 400 Gb/s. The contention point is not the network anymore, but the applications which centralize data transfers and do not exploit parallelism capacities from datacenters which include many servers (and especially many network interfaces – NIC). The PicNic project addresses this issue by allowing applications to exploit

network cards available in a datacenter, remotely, in order to optimize transfers (hence the acronym PicNic). The objective is to design a set of system services for massive data transfer between datacenters, exploiting distribution and parallelisation of networks flows.

The consortium is composed of several partners: Laboratoire d'Informatique du Parallélisme, Institut de Cancérologie de l'Ouest / Informatique, Institut de Recherche en Informatique de Toulouse, Laboratoire des Sciences du Numérique de Nantes, Laboratoire d'Informatique de Grenoble, and Nutanix France.

PiCNiC is running for 42 months (starting in Sept 2021 with an allocated budget of 495k€, 170k€ for STACK).

10.3.4 PIA 4

OTPaaS

Participants: Hélène Coullon, Remous-Aris Koutsiamanis, Adrien Lebre (*STACK representative*), Thomas Ledoux, Jean-Marc Menaud, Jacques Noyé, Mario Südholt.

The OTPaaS project targets the design and development of a complete software stack to administrate and use edge infrastructures for the industry sector. The consortium brings together national and user technology suppliers from major groups (Atos / Bull, Schneider Electric, Valeo) and SMEs / ETIs (Agileo Automation, Mydatamodels, Dupliprint, Solem, Tridimeo, Prosyst, Soben), with a strong support from major French research institutes (CEA, Inria, IMT, CAPTRONIC). The project started in October 2021 for a period of 36 months with an overall budget of 56M€ (1.2M€ for STACK).

The OTPaaS platform objectives are:

- To be built on National and sovereign technologies for the edge cloud.
- To be validated by industrial demonstrators of multisectoral use cases.
- To be followed and supported by ambitious industrialization programs.
- To be accompanied by a massive campaign to promote its use by SMEs / midcaps.
- To integrate solutions for controlling energy consumption.
- To be compliant with the Gaia-X ecosystem.

10.3.5 Etoiles Montantes

VeRDi

Participants: Emile Cadorel, Hélène Coullon (*Coordinator*), Simon Robillard.

VeRDi is an acronym for Verified Reconfiguration Driven by execution. The VeRDi project has been funded by the French region Pays De La Loire where Nantes is located. The project started in November 2018 and ended on June 2021 (extended due to Covid-19) with an allocated budget of 172800€.

It aimed at addressing distributed software reconfiguration in an efficient and verified way. The aim of the VeRDi project was to build an argued disruptive view of the problem. To do so, we wanted to validate the work already performed on the deployment in the team and extend it to reconfiguration.

10.4 Regional initiatives

10.4.1 Interregional activities

Oncoshare

Participants: Mario Südholt (*Coordinator*).

The ONCOSHARE project (ONCOlogy big data SHARing for Research) will demonstrate, through a multidisciplinary cooperation within the Western CANCEROPOLE network, the feasibility and the added value of a common cancer patient-centered information network for in-silico research. The STACK team will work on challenges to the security and the privacy of user data in this context.

This project is funded by three French regions from 2018-2021 with a global envelope of 150 K€ (and 10 K€ for our team).

10.4.2 Nantes excellency initiative in Medecine and Informatics (NExT)

SysMics

Participants: Jean-Marc Menaud, Mario Südholt (*Coordinator*).

The SysMics project aims at federating the NExT scientific community toward a common objective: anticipate the emergence of medicine systems medicine by co-developing three approaches in population-scale genomics: genotyping by sequencing, cell-by-cell profiling and microbiome analysis.

STACK investigates new means for secure and privacy-aware computations in the context of personalized medicine, notably genetic analyses. This year we have proposed a novel framework for privacy-preserving collaborative GWAS performed in the cloud. Indeed, our proposal is the first framework which combines a hybrid cloud deployment with a set of four security mechanisms: digital watermarking, homomorphic encryption, meta-data de-identification and the Intel Software Guard Extensions technology. [14]

This project's financing amounts to a global envelope of 150 K€ (10 K€ for our team) from 2018-22.

SHLARC

Participants: Sirine Sayadi, Mario Südholt (*Coordinator*).

The SHLARC project is an international network involving more than 20 partners from more than 15 countries located on four continents. The network aims at improving HLA imputation techniques in the domain of immunobiology, notably by investigation better computational methods for the corresponding biomedical analyses. The ambition of the SHLARC is to bring together international expertise to solve essential questions on immune-related pathologies through innovative algorithms and powerful computation tool development. To achieve this goal, we determined 3 main objectives • Data. By bringing together scientists from around the world, we will collectively increase the amount of SNP+HLA data available, both in terms of quantity and diversity. • Applied mathematical and computer sciences. We will further optimize SNP-HLA imputation methods using the attribute-bagging HIBAG tool, and particularly for genetically diverse and admixed populations. • Accessibility and service to the scientific community. Following the Haplotype Reference Consortium (HRC) initiative, the network envisions building a free, user-friendly webserver where researchers can access improved imputation protocols by simply uploading their data and obtaining the best possible HLA imputation for their dataset.

In this context, the STACK team is working on improved analysis techniques that harness distributed infrastructures. This year we have motivated the use of distributed implementations for FAMD analyses in the context of the development of a personalised medicine application called KITAPP. We have presented a new distribution method for FAMD and evaluated its implementation in a multi-site setting based on real data. [19]

This project is funded from 2019-22 with a global envelope of 100 K€ (and 5 K€ for our team)

11 Dissemination

11.1 Promoting scientific activities

11.1.1 Scientific events: organisation

Member of the organizing committees

- Adrien Lebre is a member of the steering committee of the international conference of Fog and Edge Computing (ICFEC).
- Mario Südholt is a member of the steering committee of the international conference <Programming>.

11.1.2 Scientific events: selection

Member of the conference program committees

- H el ene Coullon: publicity co-chair CCGrid 2021, ICCS 2021, FormaliSE 2021, SBAC-PAD 2021.
- Adrien Lebre: CCGRID 2021, ICFEC 2021, UCC 2021.
- Jean-Marc Menaud: SDS'21, SMARTGREENS'21
- Remous Aris Koutsiamanis: IEEE ICC 2021, IEEE CSCN 2021

11.1.3 Journal

Member of the editorial boards

- Adrien Lebre: Associate Editor of the IEEE Transactions on Cloud Computing.
- Mario Südholt: member of the advisory board (aka. steering committee) of The Programming Journal.

Reviewer - reviewing activities

- Thomas Ledoux has been a reviewer for the following journal: Journal of Cloud Computing (Springer)

11.1.4 Scientific expertise

- A. Lebre contributed to the Allistene working group dedicated to the definition of “Plan de relance: strat egie d’acc el eration Cloud et verdissement du num erique” roadmap, March 2021.
- A. Lebre contributed to the working group dedicated to the writing of “Livre blanc Cloud de confiance”, Think tank Digital new deal, May 2021 (further information [here](#).)
- A. Lebre contributes as the IMT representative to the working group dedicated to the organisation and structuring of the PEPR Cloud action (Budget overall: 56M ).
- A. Lebre contributed to the writing of “Livre blanc sur l’Internet des objets (IoT)”, Inria, Dec 2021.
- T. Ledoux was a reviewer for the ADEME call Perfecto 2021.
- J.-M. Menaud was an expert for the 2021 CIR evaluation

11.1.5 Research administration

- A. Lebre is a member of the executive and architect committees of the Grid'5000 GIS (Groupement d'intérêt scientifique).
- A. Lebre is a co-director of the <I/O> Lab, a joint lab between Inria and Orange Labs.
- A. Lebre is a member of the scientific committee of the joint lab between Inria and Nokia Bell Labs.
- H. Coullon is vice-president of the French ACM SIGOPS group.
- H. Coullon is co-chair of the working group YODA (trustworthY and Optimal Dynamic Adaptation) in the national reasearch group GDR GPL (software engineering and languages).
- H. Coullon is co-chair of the working group CyclOps (DevOps, deployment and reconfiguration, lifecycle etc.) of the joint laboratory between Inria and Orange (<I/O> Lab).
- J. Noyé is the deputy head of the Automation, Production and Computer Sciences Department of IMT Atlantique.
- M. Südholt is a deputy director of the graduate school MathSTIC that covers the two regions Pays de la Loire and Brittany. He is responsible for the management of doctoral studies in the Nantes region.
- M. Südholt is co-chair of the working group Security of the joint laboratory between Inria and Orange (<I/O> Lab).
- J.-M. Menaud is the organizer of "Pôle Science du Logiciel et des Systèmes Distribués" in Laboratoire des Sciences du Numérique à Nantes (LS2N).
- J.-M. Menaud is co-chair of the working group Energy of the joint laboratory between Inria and Orange (<I/O> Lab)
- J.-M. Menaud is in charge of the academic relations for the Automation, Production and Computer Sciences Department of IMT Atlantique.
- J.-M. Menaud is involved in the GIS (Groupement d'Intérêt Scientifique) VITTORIA (VirTual inTegrative Oncology Research and InnovAtion) and PERLE (Pôle d'Excellence de la Recherche Ligérienne en Energie).

11.2 Teaching - Supervision - Juries

11.2.1 Teaching

- T. Ledoux is the head of the **apprenticeship program in Software Engineering FIL**. This 3-year program leads to the award of a Master degree in Software Engineering from the IMT Atlantique.
- H. Coullon is responsible for the Computer Science domain of the new apprenticeship program in Industry 4.0 (FIT) of IMT Atlantique. This 3-year program leads to the award of a Master degree in Industry 4.0 from the IMT Atlantique.
- T. Ledoux has been the head of the Filière informatique nantaise since Sept. 2020. This entity, created by the University of Nantes, Centrale Nantes and IMT Atlantique, aims to bring together the main players in Computer Science training in Nantes to ensure a coherent and ambitious training offer that meets the present and future challenges of Computer Science. It is organized around a Council made up of representatives from the academic and socio-economic worlds.
- J. Noyé is the deputy head of the Automation, Production and Computer Sciences Department of IMT Atlantique.
- M. Südholt is the representative for MSc-level and PhD-level studies of the API department of IMT Atlantique.

11.2.2 Supervision

- PhD: Marie Delavergne, director: A. Lebre.
- PhD: David Espinel, director: A. Lebre (defended Sept. 2021).
- PhD: Geo Johns Anthony, director: A. Lebre (since Apr. 2021).
- PhD: Dimitri Saingre, advisor: T. Ledoux, director: J-M. Menaud (defended Dec. 2021).
- PhD: Jolan Philippe, advisors: H. Coullon, M. Tisi (NaoMod), director: G. Sunye (NaoMod).
- PhD: Maxime Belair, advisor: S. Laniepce (Orange Labs), director: J.-M. Menaud (defended Dec. 2021).
- PhD: Wilmer Garzon, advisor: M. Südholt.
- PhD: Sirine Sayadi, advisor: M. Südholt.
- PhD: Pierre Jacquet, advisor: T. Ledoux (since Oct. 2021).
- PhD: Hiba Awad, advisor: T. Ledoux (since Nov. 2021).
- PhD: Antoine Omond, advisor: H. Coullon, director: T. Ledoux (since Dec. 2021).
- Postdoc: Rémy Pottier, advisor: J.-M. Menaud (until Oct. 2021).
- Postdoc: Simon Robillard, advisor: H. Coullon (until June 2021).
- Engineer: Emile Cadorel, advisor: H. Coullon (until Jan. 2021).
- Postdoc: Brice Nédelec, advisor: A. Lebre (from Sept. to Dec. 2021).
- Post-doc: Abdelghani Alidra, advisor: T. Ledoux (since May 2021).
- Engineer: Ronan-Alexandre Cherrueau, advisor: A. Lebre (until July 2021).
- Engineer: Brice Nédelec, advisor: T. Ledoux (until Mar. 2021).

11.2.3 Juries

- A. Lebre was a reviewer of the Phd Committee of Alessio Diamanti, "A novel network automation architecture: from anomaly detection to dynamic reconfiguration", Conservatoire National des Arts et Metier Paris, Dec 2021.
- T. Ledoux was a member of the PhD committee of David Espinel, "Distributing connectivity management in Cloud-Edge infrastructures using SDN-based approaches", IMT Atlantique, Sep. 07, 2021.
- T. Ledoux was a reviewer of the PhD committee of Zakaria Ournani, "Software Eco-Design: Investigating and Reducing the Energy Consumption of Software", Univ. Lille, Nov. 08, 2021.
- H. Coullon was a member of the PhD committee of Tanissia DJEMAI, "Placement optimisé de services dans les architectures Fog Computing et Internet of Things sous contraintes d'énergie, de QoS et de mobilité" (French), Université de Toulouse, Feb. 03, 2021.
- H. Coullon was a member of two selection committees for associate professors at the university of Rennes and the university of Bordeaux.

11.3 Popularization

11.3.1 Articles and contents

- T. Ledoux, “La tête dans les nuages”, Atlanstic (online), Jan. 2021.
- A. Lebre contributed to the article “La double révolution du Edge Computing”, Les echos, Sept. 2021.
- T. Ledoux, ‘IMT Atlantique et Alter Way veulent inventer le circuit court de la donnée’, POC Media, Sept. 2021.

11.3.2 Education

T. Ledoux is co-leader of the Ecolog initiative. This national project is the result of the association between the Institut du Numérique Responsable and the IT sector in Nantes, which includes the University of Nantes, IMT Atlantique and Centrale Nantes. Its ambition is to create a body of training courses available in open source and dedicated to green computing.

Lately, R.-A. Koustiamanis has been taking part in the development of two MOOCs related to the Industrial IoT, one specific to IMT Atlantique and the other one as a partnership of the German-French Academy for the Industry of the Future.

12 Scientific production

12.1 Major publications

- [1] M. Chardet, H. Coullon and S. Robillard. ‘Toward Safe and Efficient Reconfiguration with Concerto’. In: *Science of Computer Programming* 203 (1st Mar. 2021), pp. 1–31. DOI: [10.1016/j.scico.2020.102582](https://doi.org/10.1016/j.scico.2020.102582). URL: <https://hal.inria.fr/hal-03103714>.
- [2] R.-A. Cherrueau, M. Delavergne, A. Van Kempen, A. Lebre, D. Pertin, J. Rojas Balderrama, A. Simonet and M. Simonin. ‘EnosLib: A Library for Experiment-Driven Research in Distributed Computing’. In: *IEEE Transactions on Parallel and Distributed Systems* (30th Sept. 2021), pp. 1–15. DOI: [10.1109/TPDS.2021.3111159](https://doi.org/10.1109/TPDS.2021.3111159). URL: <https://hal.inria.fr/hal-03324177>.
- [3] D. Espinel Sarmiento, A. Lebre, L. Nussbaum and A. Chari. ‘Decentralized SDN Control Plane for a Distributed Cloud-Edge Infrastructure: A Survey’. In: *Communications Surveys and Tutorials, IEEE Communications Society*. IEEE Communications Surveys & Tutorials 23.1 (25th Feb. 2021), pp. 256–281. DOI: [10.1109/COMST.2021.3050297](https://doi.org/10.1109/COMST.2021.3050297). URL: <https://hal.archives-ouvertes.fr/hal-03119901>.
- [4] T. Lagos Jenschke, R.-A. Koutsiamanis, G. Papadopoulos and N. Montavont. ‘ODESe: On-Demand Selection for multi-path RPL networks’. In: *Ad Hoc Networks* 114 (2021), p. 102431. DOI: [10.1016/j.adhoc.2021.102431](https://doi.org/10.1016/j.adhoc.2021.102431). URL: <https://hal.archives-ouvertes.fr/hal-03122968>.
- [5] D. Saingre, T. Ledoux and J.-M. Menaud. ‘The cost of immortality: A Time To Live for smart contracts’. In: *ISCC 2021 - 26th IEEE Symposium on Computers and Communications*. Athènes, Greece: IEEE, 5th Sept. 2021. URL: <https://hal.inria.fr/hal-03300144>.

12.2 Publications of the year

International journals

- [6] A. Bruniaux, R.-A. Koutsiamanis, G. Papadopoulos and N. Montavont. ‘Defragmenting the 6LoWPAN Fragmentation Landscape: A Performance Evaluation’. In: *Sensors* 21.5 (2nd Mar. 2021), pp. 1–20. DOI: [10.3390/s21051711](https://doi.org/10.3390/s21051711). URL: <https://hal.archives-ouvertes.fr/hal-03352772>.
- [7] M. Chardet, H. Coullon and S. Robillard. ‘Toward Safe and Efficient Reconfiguration with Concerto’. In: *Science of Computer Programming* 203 (1st Mar. 2021), pp. 1–31. DOI: [10.1016/j.scico.2020.102582](https://doi.org/10.1016/j.scico.2020.102582). URL: <https://hal.inria.fr/hal-03103714>.

- [8] R.-A. Cherrueau, M. Delavergne, A. Van Kempen, A. Lebre, D. Pertin, J. Rojas Balderrama, A. Simonet and M. Simonin. ‘EnosLib: A Library for Experiment-Driven Research in Distributed Computing’. In: *IEEE Transactions on Parallel and Distributed Systems* 33.6 (1st June 2022), pp. 1464–1477. DOI: [10.1109/TPDS.2021.3111159](https://doi.org/10.1109/TPDS.2021.3111159). URL: <https://hal.inria.fr/hal-03324177>.
- [9] D. Espinel Sarmiento, A. Lebre, L. Nussbaum and A. Chari. ‘Decentralized SDN Control Plane for a Distributed Cloud-Edge Infrastructure: A Survey’. In: *Communications Surveys and Tutorials, IEEE Communications Society. IEEE Communications Surveys & Tutorials* 23.1 (25th Feb. 2021), pp. 256–281. DOI: [10.1109/COMST.2021.3050297](https://doi.org/10.1109/COMST.2021.3050297). URL: <https://hal.archives-ouvertes.fr/hal-03119901>.
- [10] D. Hauweele, R.-A. Koutsiamanis, B. Quoitin and G. Papadopoulos. ‘Thorough Performance Evaluation & Analysis of the 6TiSCH Minimal Scheduling Function (MSF)’. In: *Journal of Signal Processing Systems* (2nd June 2021). DOI: [10.1007/s11265-021-01668-w](https://doi.org/10.1007/s11265-021-01668-w). URL: <https://hal.archives-ouvertes.fr/hal-03352745>.
- [11] T. Lagos Jenschke, R.-A. Koutsiamanis, G. Papadopoulos and N. Montavont. ‘ODESe: On-Demand Selection for multi-path RPL networks’. In: *Ad Hoc Networks* 114 (2021), p. 102431. DOI: [10.1016/j.adhoc.2021.102431](https://doi.org/10.1016/j.adhoc.2021.102431). URL: <https://hal.archives-ouvertes.fr/hal-03122968>.
- [12] D. Saingre, T. Ledoux and J.-M. Menaud. ‘Measuring performances and footprint of blockchains with BCTMark: a case study on Ethereum smart contracts energy consumption’. In: *Cluster Computing* (2021). DOI: [10.1007/s10586-021-03441-x](https://doi.org/10.1007/s10586-021-03441-x). URL: <https://hal.inria.fr/hal-03330325>.

International peer-reviewed conferences

- [13] M. Belair, S. Laniece and J.-M. Menaud. ‘SNAPPY: Programmable Kernel-Level Policies for Containers’. In: SAC 2021: 36th ACM/SIGAPP Symposium On Applied Computing. Gwangju / Virtual, South Korea, 22nd Mar. 2021, pp. 1636–1645. DOI: [10.1145/3412841.3442037](https://doi.org/10.1145/3412841.3442037). URL: <https://hal.inria.fr/hal-03108231>.
- [14] F.-z. Boujdad, D. Niyitegeka, R. Bellafqira, G. Coatrieux, E. Génin and M. Südholt. ‘A Hybrid Cloud Deployment Architecture For Privacy-preserving Collaborative Genome-Wide Association Studies’. In: ICDF2C - 12th EAI International Conference on Digital Forensics & Cyber Crime. Singapour, Singapore, 7th Dec. 2021. URL: <https://hal.inria.fr/hal-03331063>.
- [15] R.-A. Cherrueau, M. Delavergne and A. Lebre. ‘Geo-Distribute Cloud Applications at the Edge’. In: EURO-PAR 2021 - 27th International European Conference on Parallel and Distributed Computing. Lisbon, Portugal, 30th Aug. 2021, pp. 1–14. DOI: [10.1007/978-3-030-85665-6_19](https://doi.org/10.1007/978-3-030-85665-6_19). URL: <https://hal.inria.fr/hal-03212421>.
- [16] K. Manaouil and A. Lebre. ‘Kubernetes WANWide: a Deployment Scenario to Expose and Use Edge Computing Resources?’ In: PDP 2021 - 29th Euromicro International Conference on Parallel, Distributed and Network-Based Processing. Valladolid, Spain: IEEE, 10th Mar. 2021, pp. 193–197. DOI: [10.1109/PDP52278.2021.00038](https://doi.org/10.1109/PDP52278.2021.00038). URL: <https://hal.inria.fr/hal-03212370>.
- [17] J. Philippe, M. Tisi, H. Coullon and G. Sunyé. ‘Executing Certified Model Transformations on Apache Spark’. In: SLE 2021: 14th ACM SIGPLAN International Conference on Software Language Engineering. SLE 2021: Proceedings of the 14th ACM SIGPLAN International Conference on Software Language Engineering. Chicago IL, United States, 22nd Nov. 2021, pp. 36–48. DOI: [10.1145/3486608.3486901](https://doi.org/10.1145/3486608.3486901). URL: <https://hal.archives-ouvertes.fr/hal-03343942>.
- [18] D. Saingre, T. Ledoux and J.-M. Menaud. ‘The cost of immortality: A Time To Live for smart contracts’. In: ISCC 2021 - 26th IEEE Symposium on Computers and Communications. Athènes, Greece: IEEE, 5th Sept. 2021. URL: <https://hal.inria.fr/hal-03300144>.
- [19] S. Sayadi, E. Geffard, M. Südholt, N. Vince and P.-A. Gourraud. ‘Secure distribution of Factor Analysis of Mixed Data (FAMD) and its application to personalized medicine of transplanted patients’. In: AINA 2021: 35th International Conference on Advanced Information Networking and Applications. Toronto, Canada, 24th Apr. 2021. DOI: [10.1007/978-3-030-75100-5_44](https://doi.org/10.1007/978-3-030-75100-5_44). URL: <https://hal.archives-ouvertes.fr/hal-03141653>.

Conferences without proceedings

- [20] M. Delavergne, R.-A. Cherrueau and A. Lebre. ‘A service mesh for collaboration between geo-distributed services: the replication case’. In: AMP 2021 - Workshop on Agility with Microservices Programming. Online, France, 14th June 2021, pp. 1–8. URL: <https://hal.inria.fr/hal-03282425>.

Doctoral dissertations and habilitation theses

- [21] M. Bélaïr. ‘Defense against attacks with regard to new virtualisation kinds’. Ecole nationale supérieure Mines-Télécom Atlantique, 8th Dec. 2021. URL: <https://tel.archives-ouvertes.fr/tel-03520546>.
- [22] D. Espinel Sarmiento. ‘Distributing connectivity management in Cloud-Edge infrastructures using SDN-based approaches’. Ecole nationale supérieure Mines-Télécom Atlantique, 7th Sept. 2021. URL: <https://tel.archives-ouvertes.fr/tel-03394285>.
- [23] D. Saingre. ‘Understanding the energy consumption of blockchain technologies : a focus on smart contracts’. Ecole nationale supérieure Mines-Télécom Atlantique, 13th Dec. 2021. URL: <https://tel.archives-ouvertes.fr/tel-03546651>.

Reports & preprints

- [24] A. Lebre, B. Nédelec and A. Van Kempen. *AS-cast: Lock Down the Traffic of Decentralized Content Indexing at the Edge*. RR-9418. Inria Rennes, 3rd Sept. 2021, pp. 1–21. URL: <https://hal.inria.fr/hal-03333669>.

12.3 Cited publications

- [25] *Akamai Cloudlets*. (Accessed: 2018-03-08). 2018. URL: <http://cloudlets.akamai.com>.
- [26] *Amazon Lambda@Edge*. (Accessed: 2018-03-08). 2018. URL: <https://aws.amazon.com/lambda/edge/>.
- [27] C. Atkinson, T. Schulze and S. Klingert. ‘Facilitating Greener IT through Green Specifications’. In: *IEEE Software* 31.3 (May 2014), pp. 56–63. DOI: [10.1109/MS.2014.19](https://doi.org/10.1109/MS.2014.19). URL: <http://dx.doi.org/10.1109/MS.2014.19>.
- [28] F. Baude, D. Caromel, C. Dalmaso, M. Danelutto, V. Getov, L. Henrio and C. Pérez. ‘GCM: a grid extension to Fractal for autonomous distributed components’. English. In: *annals of telecommunications* 64.1-2 (2009), pp. 5–24. DOI: [10.1007/s12243-008-0068-8](https://doi.org/10.1007/s12243-008-0068-8). URL: <http://dx.doi.org/10.1007/s12243-008-0068-8>.
- [29] F. Baude, L. Henrio and C. Ruz. ‘Programming distributed and adaptable autonomous components—the GCM/ProActive framework’. In: *Software: Practice and Experience* (May 2014). URL: <https://hal.inria.fr/hal-01001043>.
- [30] M.-M. Bazm, M. Lacoste, M. Südholt and J.-M. Menaud. ‘Side-Channels Beyond the Cloud Edge : New Isolation Threats and Solutions’. In: *IEEE International Conference on Cyber Security in Networking (CSNet) 2017*. Rio de Janeiro, Brazil, Oct. 2017. URL: <https://hal.inria.fr/hal-01593144>.
- [31] G. Blair, T. Coupaye and J.-B. Stefani. ‘Component-based architecture: the Fractal initiative’. In: *Annals of telecommunications* 64.1 (Feb. 2009), pp. 1–4. DOI: [10.1007/s12243-009-0086-1](https://doi.org/10.1007/s12243-009-0086-1). URL: <https://doi.org/10.1007/s12243-009-0086-1>.
- [32] P. Bodik, R. Griffith, C. Sutton, A. Fox, M. I. Jordan and D. A. Patterson. ‘Automatic Exploration of Datacenter Performance Regimes’. In: *Proceedings of the 1st Workshop on Automated Control for Datacenters and Clouds*. ACDC ’09. Barcelona, Spain: ACM, 2009, pp. 1–6. DOI: [10.1145/1555271.1555273](https://doi.org/10.1145/1555271.1555273). URL: <http://doi.acm.org/10.1145/1555271.1555273>.

- [33] F. Bonomi, R. Milito, J. Zhu and S. Addepalli. 'Fog computing and its role in the internet of things'. In: *Proceedings of the first edition of the MCC workshop on Mobile cloud computing*. ACM. 2012, pp. 13–16.
- [34] F.-z. Boujdad and M. Südholt. 'Constructive Privacy for Shared Genetic Data'. In: *CLOSER 2018 - 8th International Conference on Cloud Computing and Services Science*. Proceedings of CLOSER 2018. Funchal, Madeira, Portugal, Mar. 2018, pp. 1–8. URL: <https://hal.archives-ouvertes.fr/hal-01692620>.
- [35] O. Bračevac, S. Erdweg, G. Salvaneschi and M. Mezini. 'CPL: A Core Language for Cloud Computing'. In: *Proceedings of the 13th International Conference on Modularity (MODULARITY 2016)*. Málaga, Spain: ACM Press, 2016, pp. 94–105. URL: <http://doi.acm.org/10.1145/2889443.2889452>.
- [36] F. Brasileiro, G. Silva, F. Araújo, M. Nóbrega, I. Silva and G. Rocha. 'Fogbow: A Middleware for the Federation of IaaS Clouds'. In: *The 16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*. IEEE. 2016, pp. 531–534.
- [37] H. Bruneliere, Z. Al-Shara, F. Alvares, J. Lejeune and T. Ledoux. 'A Model-based Architecture for Autonomic and Heterogeneous Cloud Systems'. In: *CLOSER 2018 - 8th International Conference on Cloud Computing and Services Science*. Vol. 1. Best Paper Award. Funchal, Portugal, Mar. 2018, pp. 201–212. DOI: [10.5220/0006773002010212](https://doi.org/10.5220/0006773002010212). URL: <https://hal.archives-ouvertes.fr/hal-01705248>.
- [38] E. Bruneton, T. Coupaye, M. Leclercq, V. Quema and J.-B. Stefani. 'An Open Component Model and Its Support in Java'. In: *Component-Based Software Engineering*. Ed. by I. Crnkovic, J. A. Stafford, H. W. Schmidt and K. Wallnau. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 7–22.
- [39] P. Carbone, A. Katsifodimos, S. Ewen, V. Markl, S. Haridi and K. Tzoumas. 'Apache Flink?: Stream and Batch Processing in a Single Engine'. In: *IEEE Data Eng. Bull.* 38 (2015), pp. 28–38.
- [40] V. Cardellini, F. L. Presti, M. Nardelli and G. R. Russo. 'Towards Hierarchical Autonomous Control for Elastic Data Stream Processing in the Fog'. In: *Euro-Par 2017: Parallel Processing Workshops: Euro-Par 2017 International Workshops, Santiago de Compostela, Spain, August 28-29, 2017, Revised Selected Papers*. Vol. 10659. Springer. 2018, pp. 106–117.
- [41] F.-C. Chang, H.-C. Huang and H.-M. Hang. 'Combined Encryption and Watermarking Approaches for Scalable Multimedia Coding'. In: *Advances in Multimedia Information Processing - PCM 2004: 5th Pacific Rim Conference on Multimedia, Tokyo, Japan, November 30 - December 3, 2004. Proceedings, Part III*. Ed. by K. Aizawa, Y. Nakamura and S. Satoh. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 356–363.
- [42] R.-A. Cherrueau, R. Douence and M. Südholt. 'A Language for the Composition of Privacy-Enforcement Techniques'. In: *IEEE RATSP 2015, The 2015 IEEE International Symposium on Recent Advances of Trust, Security and Privacy in Computing and Communications*. Helsinki, Finland, Aug. 2015. URL: <https://hal.inria.fr/hal-01168303>.
- [43] R.-A. Cherrueau, A. Lebre, D. Pertin, F. Wuhib and J. Soares. 'Edge Computing Resource Management System: a Critical Building Block! Initiating the debate via OpenStack'. In: *The USENIX Workshop on Hot Topics in Edge Computing (HotEdge'18)*. July 2018.
- [44] H. E. Chihoub, S. Ibrahim, Y. Li, G. Antoniu, M. S. Perez and L. Bouge. 'Exploring Energy-Consistency Trade-Offs in Cassandra Cloud Storage System'. In: *27th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*. Oct. 2015, pp. 146–153.
- [45] B. Confais, A. Lebre and B. Parrein. 'An Object Store Service for a Fog/Edge Computing Infrastructure based on IPFS and Scale-out NAS'. In: *1st IEEE International Conference on Fog and Edge Computing-ICFEC'2017*. 2017.
- [46] J. C. Corbett, J. Dean, M. Epstein, A. Fikes, C. Frost, J. J. Furman, S. Ghemawat, A. Gubarev, C. Heiser and P. Hochschild. 'Spanner: Google's globally distributed database'. In: *ACM Transactions on Computer Systems (TOCS)* 31.3 (2013), p. 8.

- [47] R. D. Cosmo, J. Mauro, S. Zacchiroli and G. Zavattaro. 'Aeolus: A component model for the cloud'. In: *Information and Computation* 239.Supplement C (2014), pp. 100–121. DOI: <https://doi.org/10.1016/j.ic.2014.11.002>. URL: <http://www.sciencedirect.com/science/article/pii/S0890540114001424>.
- [48] H. Coullon, J. Bigot and C. Pérez. 'Extensibility and Composability of a Multi-Stencil Domain Specific Framework'. In: *International Journal of Parallel Programming* (Nov. 2017). DOI: [10.1007/s10766-017-0539-5](https://doi.org/10.1007/s10766-017-0539-5). URL: <https://doi.org/10.1007/s10766-017-0539-5>.
- [49] H. Coullon, G. Le Louët and J.-M. Menaud. 'Virtual Machine Placement for Hybrid Cloud using Constraint Programming'. In: *ICPADS 2017*. 2017.
- [50] H. Coullon, D. Pertin and C. Pérez. 'Production Deployment Tools for IaaS: an Overall Model and Survey'. In: *IEEE International Conference on Future Internet of Things and Cloud (FiCloud) 2017*. Prague, Czech Republic, Aug. 2017. URL: <https://hal.inria.fr/hal-01532489>.
- [51] I. Cuadrado-Cordero, A.-C. Orgerie and J.-M. Menaud. 'Comparative Experimental Analysis of the Quality-of-Service and Energy-Efficiency of VMs and Containers' Consolidation for Cloud Applications'. In: *SoftCOM: International Conference on Software, Telecommunications and Computer Networks*. 2017.
- [52] P.-C. David, T. Ledoux, M. Léger and T. Coupaye. 'FPath and FScript: Language support for navigation and reliable reconfiguration of Fractal architectures'. In: *annals of telecommunications - annales des télécommunications* 64.1 (Feb. 2009), pp. 45–63. DOI: [10.1007/s12243-008-0073-y](https://doi.org/10.1007/s12243-008-0073-y). URL: <https://doi.org/10.1007/s12243-008-0073-y>.
- [53] F. A. De Oliveira, T. Ledoux and R. Sharrock. 'A framework for the coordination of multiple autonomic managers in cloud environments'. In: *Self-Adaptive and Self-Organizing Systems (SASO), 2013 IEEE 7th International Conference on*. IEEE. 2013, pp. 179–188.
- [54] W. Felter, A. Ferreira, R. Rajamony and J. Rubio. 'An updated performance comparison of virtual machines and linux containers'. In: *Performance Analysis of Systems and Software (ISPASS), 2015 IEEE International Symposium On*. IEEE. 2015, pp. 171–172.
- [55] A. Flissi, J. Dubus, N. Dolet and P. Merle. 'Deploying on the Grid with DeployWare'. In: *Eighth IEEE International Symposium on Cluster Computing and the Grid*. France, May 2008, pp. 177–184. DOI: [10.1109/CCGRID.2008.59](https://hal.archives-ouvertes.fr/hal-00259836). URL: <https://hal.archives-ouvertes.fr/hal-00259836>.
- [56] P. Garcia Lopez, A. Montresor, D. Epema, A. Datta, T. Higashino, A. Iamnitchi, M. Barcellos, P. Felber and E. Riviere. 'Edge-centric Computing: Vision and Challenges'. In: *SIGCOMM Comput. Commun. Rev.* 45.5 (Sept. 2015), pp. 37–42. DOI: [10.1145/2831347.2831354](https://doi.org/10.1145/2831347.2831354). URL: <http://doi.acm.org/10.1145/2831347.2831354>.
- [57] Í. Goiri, W. Katsak, K. Le, T. D. Nguyen and R. Bianchini. 'Parasol and GreenSwitch: Managing Datacenters Powered by Renewable Energy'. In: *SIGARCH Comput. Archit. News* 41.1 (Mar. 2013), pp. 51–64.
- [58] Í. Goiri, K. Le, T. D. Nguyen, J. Guitart, J. Torres and R. Bianchini. 'GreenHadoop: Leveraging Green Energy in Data-processing Frameworks'. In: *Proceedings of the 7th ACM European Conference on Computer Systems*. EuroSys '12. Bern, Switzerland: ACM, 2012, pp. 57–70. DOI: [10.1145/2168836.2168843](https://doi.org/10.1145/2168836.2168843). URL: <http://doi.acm.org/10.1145/2168836.2168843>.
- [59] H. Goudarzi and M. Pedram. 'Geographical Load Balancing for Online Service Applications in Distributed Datacenters'. In: *in IEEE international conference on cloud computing (CLOUD 2013)*. 2013.
- [60] L. Gu, D. Zeng, A. Barnawi, S. Guo and I. Stojmenovic. 'Optimal Task Placement with QoS Constraints in Geo-Distributed Data Centers Using DVFS'. In: *IEEE Transactions on Computers* 64.7 (July 2015), pp. 2049–2059. DOI: [10.1109/TC.2014.2349510](https://doi.org/10.1109/TC.2014.2349510). URL: <http://dx.doi.org/10.1109/TC.2014.2349510>.
- [61] K. Ha, Y. Abe, T. Eiszler, Z. Chen, W. Hu, B. Amos, R. Upadhyaya, P. Pillai and M. Satyanarayanan. 'You Can Teach Elephants to Dance: Agile VM Handoff for Edge Computing'. In: *Proceedings of the Second ACM/IEEE Symposium on Edge Computing*. SEC '17. San Jose, California: ACM, 2017. DOI: [10.1145/3132211.3134453](https://doi.org/10.1145/3132211.3134453). URL: <http://doi.acm.org/10.1145/3132211.3134453>.

- [62] B. Han, V. Gopalakrishnan, L. Ji and S. Lee. 'Network function virtualization: Challenges and opportunities for innovations'. In: *IEEE Communications Magazine* 53.2 (Feb. 2015), pp. 90–97. DOI: [10.1109/MCOM.2015.7045396](https://doi.org/10.1109/MCOM.2015.7045396). URL: <http://dx.doi.org/10.1109/MCOM.2015.7045396>.
- [63] M. S. Hasan, F. Alvares, T. Ledoux and J. L. Pazat. 'Investigating Energy Consumption and Performance Trade-Off for Interactive Cloud Application'. In: *IEEE Transactions on Sustainable Computing* 2.2 (Apr. 2017), pp. 113–126. DOI: [10.1109/TSUSC.2017.2714959](https://doi.org/10.1109/TSUSC.2017.2714959). URL: <http://dx.doi.org/10.1109/TSUSC.2017.2714959>.
- [64] F. Hermenier, J. Lawall and G. Muller. 'Btrplace: A flexible consolidation manager for highly available applications'. In: *IEEE Transactions on dependable and Secure Computing* 10.5 (2013), pp. 273–286.
- [65] F. Hermenier, A. Lebre and J.-M. Menaud. 'Cluster-wide Context Switch of Virtualized Jobs'. In: *Proceedings of the Virtualization Technologies in Distributed Computing Workshop (co-located with ACM HPDC'10)*. HPDC '10. Chicago, Illinois: ACM, 2010, pp. 658–666. DOI: [10.1145/1851476.1851574](https://doi.org/10.1145/1851476.1851574). URL: <http://doi.acm.org/10.1145/1851476.1851574>.
- [66] F. Hermenier, X. Lorca, J.-M. Menaud, G. Muller and J. Lawall. 'Entropy: a consolidation manager for clusters'. In: *Proceedings of the 2009 ACM SIGPLAN/SIGOPS international conference on Virtual execution environments*. ACM, 2009, pp. 41–50.
- [67] M. Hirzel, S. Schneider and B. Gedik. 'SPL: An Extensible Language for Distributed Stream Processing'. In: *toplas* 39.1 (Mar. 2017). URL: <http://doi.acm.org/10.1145/3039207>.
- [68] S. Ibrahim, B. He and H. Jin. 'Towards Pay-As-You-Consume Cloud Computing'. In: *2011 IEEE International Conference on Services Computing*. July 2011, pp. 370–377. DOI: [10.1109/SCC.2011.38](https://doi.org/10.1109/SCC.2011.38). URL: <http://dx.doi.org/10.1109/SCC.2011.38>.
- [69] S. Iftikhar, S. Khan, Z. Anwar et al. 'GenInfoGuard: A Robust and Distortion-Free Watermarking Technique for Genetic Data'. In: *PLOS ONE* 10.2 (Feb. 2015), pp. 1–22. DOI: [10.1371/journal.pone.0117717](https://doi.org/10.1371/journal.pone.0117717). URL: <https://doi.org/10.1371/journal.pone.0117717>.
- [70] B. Jennings and R. Stadler. 'Resource Management in Clouds: Survey and Research Challenges'. In: *Journal of Network and Systems Management* 23.3 (July 2015), pp. 567–619. DOI: [10.1007/s10922-014-9307-7](https://doi.org/10.1007/s10922-014-9307-7). URL: <https://doi.org/10.1007/s10922-014-9307-7>.
- [71] E. Jonardi, M. A. Oxley, S. Pasricha, A. A. Maciejewski and H. J. Siegel. 'Energy cost optimization for geographically distributed heterogeneous data centers'. In: *2015 Sixth International Green and Sustainable Computing Conference (IGSC)*. Dec. 2015, pp. 1–6. DOI: [10.1109/IGCC.2015.7393677](https://doi.org/10.1109/IGCC.2015.7393677). URL: <http://dx.doi.org/10.1109/IGCC.2015.7393677>.
- [72] H. Karau, A. Konwinski, P. Wendell and M. Zaharia. *Learning Spark*. O'Reilly Media, Feb. 2015.
- [73] J. Kramer and J. Magee. 'The evolving philosophers problem: dynamic change management'. In: *IEEE Transactions on Software Engineering* 16.11 (Nov. 1990), pp. 1293–1306. DOI: [10.1109/32.60317](https://doi.org/10.1109/32.60317). URL: <http://dx.doi.org/10.1109/32.60317>.
- [74] A. Lebre, J. Pastor, A. Simonet and F. Desprez. 'Revising OpenStack to Operate Fog/Edge Computing Infrastructures'. In: *The IEEE International Conference on Cloud Engineering (IC2E)*. Apr. 2017, pp. 138–148. DOI: [10.1109/IC2E.2017.35](https://doi.org/10.1109/IC2E.2017.35). URL: <http://dx.doi.org/10.1109/IC2E.2017.35>.
- [75] G. Malewicz, M. H. Austern, A. J. Bik, J. C. Dehnert, I. Horn, N. Leiser and G. Czajkowski. 'Pregel: A System for Large-scale Graph Processing'. In: *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*. SIGMOD '10. Indianapolis, Indiana, USA: ACM, 2010, pp. 135–146. URL: <http://doi.acm.org/10.1145/1807167.1807184>.
- [76] T. L. Nguyen and A. Lebre. 'Virtual Machine Boot Time Model'. In: *Parallel, Distributed and Network-based Processing (PDP), 2017 25th Euromicro International Conference on*. IEEE, 2017, pp. 430–437.
- [77] K. Okamura and Y. Oyama. 'Load-based covert channels between Xen virtual machines'. In: *Proceedings of the 2010 ACM Symposium on Applied Computing*. ACM, 2010, pp. 173–180.

- [78] P. Oreizy, N. Medvidovic and R. N. Taylor. ‘Runtime Software Adaptation: Framework, Approaches, and Styles’. In: *Companion of the 30th International Conference on Software Engineering*. ICSE Companion ’08. Leipzig, Germany: ACM, 2008, pp. 899–910. DOI: [10.1145/1370175.1370181](https://doi.org/10.1145/1370175.1370181). URL: <http://doi.acm.org/10.1145/1370175.1370181>.
- [79] T. D. Phan, S. Ibrahim, G. Antoniu and L. Bouge. ‘On Understanding the Energy Impact of Speculative Execution in Hadoop’. In: *2015 IEEE International Conference on Data Science and Data Intensive Systems*. Dec. 2015, pp. 396–403.
- [80] F. Quesnel, A. Lebre and M. Südholt. ‘Cooperative and reactive scheduling in large-scale virtualized platforms with DVMS’. In: *Concurrency and Computation: Practice and Experience* 25.12 (2013), pp. 1643–1655.
- [81] D. Sabella, A. Vaillant, P. Kuure, U. Rauschenbach and F. Giust. ‘Mobile-Edge Computing Architecture: The role of MEC in the Internet of Things’. In: *IEEE Consumer Electronics Magazine* 5.4 (Oct. 2016), pp. 84–91. DOI: [10.1109/MCE.2016.2590118](https://doi.org/10.1109/MCE.2016.2590118). URL: <http://dx.doi.org/10.1109/MCE.2016.2590118>.
- [82] D. Serrano, S. Bouchenak, Y. Kouki, F. A. de Oliveira Jr., T. Ledoux, J. Lejeune, J. Sopena, L. Arantes and P. Sens. ‘SLA guarantees for cloud services’. In: *Future Generation Computer Systems* 54. Supplement C (2016), pp. 233–246. DOI: <https://doi.org/10.1016/j.future.2015.03.018>. URL: <http://www.sciencedirect.com/science/article/pii/S0167739X15000801>.
- [83] Q. Shen, X. Liang, X. S. Shen, X. Lin and H. Y. Luo. ‘Exploiting Geo-Distributed Clouds for a E-Health Monitoring System With Minimum Service Delay and Privacy Preservation’. In: *IEEE Journal of Biomedical and Health Informatics* 18.2 (Mar. 2014), pp. 430–439. DOI: [10.1109/JBHI.2013.2292829](https://doi.org/10.1109/JBHI.2013.2292829). URL: <http://dx.doi.org/10.1109/JBHI.2013.2292829>.
- [84] W. Shi and T. F. Wenisch. ‘Energy-Efficient Data Centers’. In: *IEEE Internet Computing* 21.4 (2017), pp. 6–7. DOI: [10.1109/MIC.2017.2911429](https://doi.org/10.1109/MIC.2017.2911429). URL: <http://dx.doi.org/10.1109/MIC.2017.2911429>.
- [85] J.-B. Stefani. ‘Components as Location Graphs’. In: *Formal Aspects of Component Software - 11th International Symposium, FACS 2014, Bertinoro, Italy, September 10-12, 2014, Revised Selected Papers*. 2014, pp. 3–23. URL: https://doi.org/10.1007/978-3-319-15317-9_1.
- [86] C. Szyperski. *Component Software: Beyond Object-Oriented Programming*. 2nd. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 2002.
- [87] Y. Taleb, S. Ibrahim, G. Antoniu and T. Cortes. ‘Characterizing Performance and Energy-Efficiency of the RAMCloud Storage System’. In: *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. June 2017, pp. 1488–1498. DOI: [10.1109/ICDCS.2017.51](https://doi.org/10.1109/ICDCS.2017.51). URL: <http://dx.doi.org/10.1109/ICDCS.2017.51>.
- [88] The Linux Foundation. *Open Network Automation Platform*. (Accessed: 2018-03-08). 2018. URL: <https://www.onap.org/>.
- [89] The OpenStack Foundation. *Cloud Edge Computing: Beyond the Data Center (White Paper)*. (Accessed: 2020-02-08). Jan. 2018. URL: <https://www.openstack.org/assets/edge/OpenStack-EdgeWhitepaper-v3-online.pdf>.
- [90] The OpenStack Foundation. *Edge Computing: Next Steps in Architecture, Design and Testing*. (Accessed: 2020-02-08). Jan. 2020. URL: <https://www.openstack.org/edge-computing/edge-computing-next-steps-in-architecture-design-and-testing>.
- [91] A. Thusoo, Z. Shao, S. Anthony, D. Borthakur, N. Jain, J. Sen Sarma, R. Murthy and H. Liu. ‘Data Warehousing and Analytics Infrastructure at Facebook’. In: *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*. SIGMOD ’10. Indianapolis, Indiana, USA: ACM Press, 2010, pp. 1013–1020. URL: <http://doi.acm.org/10.1145/1807167.1807278>.
- [92] T. White. *Hadoop: The Definitive Guide*. 4th. O’Reilly Media, Apr. 2015.
- [93] A. Wierman, Z. Liu, I. Liu and H. Mohsenian-Rad. ‘Opportunities and challenges for data center demand response’. In: *International Green Computing Conference*. Nov. 2014, pp. 1–10. DOI: [10.1109/IGCC.2014.7039172](https://doi.org/10.1109/IGCC.2014.7039172). URL: <http://dx.doi.org/10.1109/IGCC.2014.7039172>.

- [94] J. Xiao, Z. Xu, H. Huang and H. Wang. 'Security implications of memory deduplication in a virtualized environment'. In: *2013 43rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2013, pp. 1–12.
- [95] F. Xu, F. Liu, H. Jin and A. V. Vasilakos. 'Managing performance overhead of virtual machines in cloud computing: A survey, state of the art, and future directions'. In: *Proceedings of the IEEE* 102.1 (2014), pp. 11–31.
- [96] O. Yildiz, M. Dorier, S. Ibrahim, R. Ross and G. Antoniu. 'On the Root Causes of Cross-Application I/O Interference in HPC Storage Systems'. In: *2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. May 2016, pp. 750–759. DOI: [10.1109/IPDPS.2016.50](https://doi.org/10.1109/IPDPS.2016.50). URL: <http://dx.doi.org/10.1109/IPDPS.2016.50>.
- [97] O. Yildiz, A. C. Zhou and S. Ibrahim. 'Eley: On the Effectiveness of Burst Buffers for Big Data Processing in HPC Systems'. In: *2017 IEEE International Conference on Cluster Computing (CLUSTER)*. Sept. 2017, pp. 87–91. DOI: [10.1109/CLUSTER.2017.73](https://doi.org/10.1109/CLUSTER.2017.73). URL: <http://dx.doi.org/10.1109/CLUSTER.2017.73>.
- [98] M. Zaharia, T. Das, H. Li, T. Hunter, S. Shenker and I. Stoica. 'Discretized Streams: Fault-tolerant Streaming Computation at Scale'. In: *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*. SOSP '13. Farmington, Pennsylvania: ACM, 2013, pp. 423–438. DOI: [10.1145/2517349.2522737](https://doi.org/10.1145/2517349.2522737). URL: <http://doi.acm.org/10.1145/2517349.2522737>.